arXiv:1412.0753v4 [stat.ME] 23 Feb 2016

Consistent clustering using an ℓ_1 fusion penalty

PETER RADCHENKO AND GOURAB MUKHERJEE

Abstract

We study the large sample behavior of a convex clustering framework, which minimizes the sample within cluster sum of squares under an ℓ_1 fusion constraint on the cluster centroids. This recently proposed approach has been gaining in popularity, however, its asymptotic properties have remained mostly unknown. Our analysis is based on a novel representation of the sample clustering procedure as a sequence of cluster splits determined by a sequence of maximization problems. We use this representation to provide a simple and intuitive formulation for the population clustering procedure, and demonstrate that the sample procedure consistently estimates its population analog. We also derive the corresponding rates of convergence. The proof conducts a careful simultaneous analysis of a collection of M-estimation problems, whose cardinality grows together with the sample size.

Based on the new perspectives gained from the asymptotic investigation, we propose a key post-processing modification of the original clustering approach. Using simulated data, we compare the proposed method with existing number of clusters and modality assessment approaches, and obtain encouraging results. We also demonstrate the applicability of our clustering method for the detection of cellular subpopulations in a single-cell virology study.

Some key words: Convex Clustering; Fusion Penalties; Number of Clusters; M-estimation; Empirical Process; Consistency; Rates of Convergence

1 Introduction

Clustering is one of the most popular statistical techniques for unsupervised classification and taxonomy detection (Hartigan, 1975; Kaufman and Rousseeuw, 2009). One serious

limitation of the traditional methods, such as k-means, is the non-convexity of the corresponding optimization problems. However, several convex clustering algorithms have been proposed in recent years (Xu *et al.*, 2004; Bach and Harchaoui, 2008; Chi and Lange, 2013). Speed and scalability of these algorithms make them increasingly popular for cluster analysis of modern massive datasets. These approaches are based on convex relaxations of the traditional non-convex clustering criteria, however, they do not naturally inherit the statistical consistency properties associated with the traditional methods. Here we study the large sample behavior of one such popular convex clustering framework, which is based on an ℓ_1 fusion penalty (Hocking *et al.*, 2011).

Consider the problem of clustering n observations, $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n$, which are sampled from a Euclidean space, $\mathbb{R}^p$. The well-studied k-means approach (MacQueen *et al.*, 1967; Hartigan, 1978; Pollard, 1981, 1982; Jain, 2010) is based on minimizing the within cluster sum of squares, $\sum_{i=1}^{n} \|\boldsymbol{x}_i - \boldsymbol{\alpha}_i\|_2^2$, with respect to the cluster centroids, $\boldsymbol{\alpha}_1, \ldots, \boldsymbol{\alpha}_n$, under the restriction that the number of distinct cluster centroids is at most k. This restriction can be viewed as an ℓ_0 constraint on the centroids. Motivated by the Lasso and its variants (Tibshirani, 1996; Tibshirani *et al.*, 2005), which successfully use the ℓ_1 constraint as a surrogate for the NP-hard ℓ_0 constraint, Hocking *et al.* (2011) consider the following modification of the k-means clustering criterion:

$$\min_{\boldsymbol{\alpha}_1,\ldots,\boldsymbol{\alpha}_n} \sum_{i=1}^{n} \|\boldsymbol{x}_i - \boldsymbol{\alpha}_i\|_2^2 \text{ subject to } \sum_{1\le i<j\le n} \left\|\boldsymbol{\alpha_i} - \boldsymbol{\alpha_j}\right\|_1 \le t. \tag{1}$$

When $t = 0$, the ℓ_1 penalty fuses all the cluster centroids together. Thus, all the observations are placed in the same cluster. When $t \ge \sum_{i<j} \|\boldsymbol{x}_i - \boldsymbol{x}_j\|_1$, we have $\boldsymbol{\alpha}_i = \boldsymbol{x}_i$ for all i, and thus, each observation forms its own cluster. Varying t between the two extremes creates a path of solutions to the regularized clustering problem. Note that the Lagrangian form of the above criterion, $\min_{\boldsymbol{\alpha}_i} \sum_{i=1}^{n} \|\boldsymbol{x}_i - \boldsymbol{\alpha}_i\|_2^2 + \lambda \sum_{i<j} \left\|\boldsymbol{\alpha_i} - \boldsymbol{\alpha_j}\right\|_1$, is separable across dimensions. Consequently, the corresponding optimization problem reduces to independently minimizing p univariate convex clustering criteria.

Thus, to understand the large sample behaviour of the multivariate solution, we only

need to focus on the analysis of the univariate clustering criterion,

$$\min_{\alpha_1,\ldots,\alpha_n} \sum_{i=1}^{n} (x_i - \alpha_i)^2 + \lambda \sum_{1 \le i < j \le n} \big|\alpha_i - \alpha_j\big|. \tag{2}$$

As the penalty parameter λ varies from 0 to ∞, each corresponding solution determines a cluster partition. We are interested in the asymptotics of the entire collection of such partitions, which we view as the outcome of the *sample clustering procedure*.

Summary of the Main Contributions. We analyze the large sample behavior of the sample clustering procedure determined by the solution path for criterion (2). We develop a simple and intuitive formulation for the population clustering procedure, and show that under some very mild regularity conditions the sample procedure consistently estimates its population analog. More specifically, we first demonstrate that the path of solutions to (2) determines a clustering tree, which can be formed by either successive merges of clusters, in a bottom up fashion, or successive splits, in a top down approach. We then study the asymptotic behavior of the full clustering tree by representing each split as a solution to a maximization problem. We define the corresponding population clustering procedure in a similar fashion, but replace sample averages with the corresponding expected values. The asymptotic analysis is significantly complicated by the fact that, unlike in the standard M-estimation setup (e.g. van der Vaart and Wellner 1996; van der Vaart 1998), the number of maximization problems at the sample level tends to infinity together with n, and the number of the corresponding population problems is infinite. We demonstrate that the sample criterion functions are uniformly close to their population counterparts. Then, a careful analysis of the continuity properties of the population solutions allows us to establish consistency of the sample clustering procedure. Furthermore, we strengthen our asymptotic results by deriving the corresponding rates of convergence.

Motivated by the results of our large sample investigation, we introduce a key post-processing modification to the sample clustering procedure. We show that the resulting procedure can be successfully used to estimate the number of clusters in the population and compare it with existing number of clusters and modality assessment approaches. Our results provide strong support for the use of ℓ_1 fusion penalties in clustering.

Connections to Related Work. In recent years, modified versions of optimization problem (1) have been studied (Hocking *et al.*, 2011; Chi and Lange, 2013; Tan and Wit-

ten, 2015). The modifications include using ℓ_2 or ℓ_∞ regularization, as well as incorporating weights (Pelckmans *et al.*, 2005; Lindsten *et al.*, 2011; Zhu *et al.*, 2014). The large sample analysis in the aforementioned papers focusses on showing that if the distance between clusters grows at a sufficiently fast rate, then the corresponding method can separate the groups perfectly. Here we consider a completely different perspective and analyze the asymptotics of the clustering approach in the classical sense of Pollard (1981). We study the clustering procedure that is applied to a random sample, and show that it consistently recovers the outcome of corresponding population procedure, which is based on the underlying probability distribution.

The criterion in (1) can be viewed (Hocking *et al.*, 2011) as a convex relaxation of the hierarchical clustering criterion (Hartigan, 1975). However, as clustering is a very mature subject, approaches built on several other philosophies are also widely used in practice. A detailed review of clustering methods can be found in Kaufman and Rousseeuw (2009). One of the most popular methods is the k-means algorithm (MacQueen *et al.*, 1967), which follows a partitioning approach for making clusters. Other popular partitioning methods, such as PAM (Kaufman and Rousseeuw, 1990) and CLARA (Kaufman and Rousseeuw, 1986), are based on the k-medoids algorithm. Density driven approaches, which provide an effective and flexible framework for clustering, include mixture model based methods, such as Fraley and Raftery (2002) and Li (2005), as well as non-parametric methods (see Li *et al.* 2007 and the references therein). In our empirical analysis we compare the performance of the proposed approach with all of the aforementioned clustering methods.

The ℓ_1 penalty, which is extensively used for variable selection (Tibshirani, 2011), also finds its use in trend filtering (Tibshirani, 2013) and high-dimensional clustering problems (Soltanolkotabi and Candés, 2012; Witten and Tibshirani, 2010). Another related approach, the fused Lasso (Rinaldo *et al.*, 2009; Tibshirani and Walther, 2005; Hoefling, 2010), deals with applications having ordered features and checks for local constancy of their associated coefficients. This approach penalizes the successive differences of the coefficients. Shen and Huang (2010); Shen *et al.* (2012); Ke *et al.* (2013); Bondell and Reich (2008) have proposed methods based on fusion penalties, which apply to all the pairwise differences of coefficients. These approaches can successfully recover the grouping structure of predictors in a high-dimensional regression setup. However, the theory developed for these methods focusses on the homogeneity of regression coefficients and cannot be applied in

the unsupervised clustering setup considered in this paper.

Organization of the Paper. In Section 2 we derive two equivalent algorithmic representations of the sample clustering procedure, which we use to formulate the corresponding population procedure. Section 3 contains our main result, which shows that, under very mild regularity conditions, the sample procedure consistently estimates its population counterpart. We also derive the corresponding rates of convergence. Our asymptotic analysis reveals that an overwhelming majority of the sample clusters are in some sense negligible. Motivated by this observation, we introduce a key post-processing modification to the clustering procedure. In Section 4 we conduct a detailed empirical analysis of our approach. More specifically, we use simulated data to show its strong performance relative to popular existing approaches for assessing modality and estimating the number of clusters. We also illustrate the use of our method in analysis of single-cell virology datasets. All the proofs, together with additional technical details, are relegated to the Supplementary Material.

2 Sample and Population Clustering Procedures

In this section we derive two equivalent representations of the sample clustering procedure. First, we develop a computationally efficient merging algorithm for producing a path of solutions to the clustering criterion (2). Then, in order to understand the large sample behavior of the solution path, we introduce an equivalent splitting procedure, which can recover all the corresponding cluster splits by solving a sequence of maximization problems. We use the splitting representation to define the population clustering procedure, and describe its basic properties.

2.1 Equivalent Representations for the Sample Solution Path

Note that a solution path for problem (2) could be produced using the highly general fused lasso algorithm in Hoefling (2010). Instead, we obtain a very simple and computationally efficient fitting procedure by analyzing our clustering criterion, (2), directly. The path algorithm we describe here is a bottom up procedure, which starts at $\lambda = 0$, with each observation forming its own cluster, and then gradually merges suitable clusters as λ increases. Fix λ, and suppose that C is one of the clusters identified by the solution to the optimization problem (2). Write α_C for the centroid of cluster C, and denote the corre-

sponding cluster average by $\overline{X}_C$. As pointed out in Hocking *et al.* (2011), the first order conditions for criterion (2) imply

$$\alpha_C = \overline{X}_C + \lambda \sum_{j,\alpha_j \neq \alpha_C} \text{sign}(\alpha_j - \alpha_C). \tag{3}$$

Note that until the cluster partition or the ordering or the centroids are modified, parameter λ is the only component on the right-hand side of the equation that can change. Thus, equation (3) provides a simple way of tracking the piecewise linear paths of the centroids α_i. Another consequence of the first order conditions is the fact that as λ increases, the only way the clusters get modified is that some of them get merged together (Hocking *et al.*, 2011). Hence, we can store the full cluster partition path by keeping track of the merges and the corresponding values of the tuning parameter λ. Algorithm 1 makes this idea precise, and Theorem 1 provides a rigorous justification.

INITIALIZE:

Sort data in ascending order and store them as $\boldsymbol{x}_n = \{x_1, \ldots, x_n\}$.

Set K, the number of clusters, equal to n. For each i in $1, ..., n$, set $C_i = \{x_i\}$.

REPEAT:

Find the adjacent centroid distances standardized by cluster sizes:

$d(j, j+1) \leftarrow \left(\overline{X}_{C_{j+1}} - \overline{X}_{C_j}\right) / (|C_j| + |C_{j+1}|)$.

Find the clusters that minimize this distance: $j^* \leftarrow \arg\min_j d(j, j+1)$.

Merge the clusters that were found: $C_{j^*} \leftarrow C_{j^*} \cup C_{j^*+1}$.

Store the above merge and the corresponding λ value: $\lambda = d(j^*, j^*+1)$.

Relabel the remaining clusters: for $j > j^*$ set $C_j \leftarrow C_{j+1}$.

Reduce the total number of clusters: $K \leftarrow K - 1$.

UNTIL $K = 1$.

OUTPUT: Sequence of cluster merges and corresponding λ values.

Algorithm 1: Merging Algorithm

The following result shows that Algorithm 1 reproduces the sequence of cluster partitions and the corresponding λ values from the optimization problem (2). In the proof, which is provided in the Supplementary Material, we also verify that the sequence of λ values, corresponding to successive merges in Algorithm 1, is increasing.

Proposition 1. *Suppose that the observations are generated from a continuous distribution. Then, with probability one, the sequence of merges and λ values produced by the merging algorithm is the same as the sequence corresponding to the optimization criterion* (2).

For the asymptotic analysis, it is helpful to recover the sequence of cluster partitions in a top down approach: we start with everything in one cluster and then split the clusters iteratively. We call a representation of the cluster C as $C = C_1 \cup C_2$ a split if $\max C_1 < \min C_2$. The full sequence of splits corresponding to the optimization problem (2) is given by the splitting procedure, described in Algorithm 2 below. Proposition 2, proved in the Supplementary Material, provides theoretical justification. In particular, it shows that each of the cluster splits is chosen to maximize the distance between the two sub-cluster means.

INITIALIZE:

Sort data in ascending order and store them as $\boldsymbol{x}_n = \{x_1, \ldots, x_n\}$.

Set the current partition of $\boldsymbol{x}_n$ to $\boldsymbol{x}_n$.

REPEAT:

Select one cluster, C, with $|C| > 1$, from a current cluster partition of $\boldsymbol{x}_n$.

Find a split partition $C = C_1 \cup C_2$, that maximizes the distance $\overline{X}_{C_2} - \overline{X}_{C_1}$.

Store the split $C = C_1 \cup C_2$ and the corresponding value $\lambda = \left(\overline{X}_{C_2} - \overline{X}_{C_1}\right)/|C|$.

Replace C with $C_1 \cup C_2$ in the current partition of $\boldsymbol{x}_n$.

UNTIL: All the clusters in the current partition of $\boldsymbol{x}_n$ are of size one.

OUTPUT: Sequence of cluster splits and corresponding λ values.

Algorithm 2: Splitting Procedure

Proposition 2. *Suppose that the observations are generated from a continuous distribution. Then, with probability one, the sequence of splits and corresponding λ values produced by the splitting procedure in Algorithm 2 exactly matches the sequence of merges and corresponding λ values produced by the merging algorithm .*

Note that, unlike the merging algorithm, the splitting procedure does not provide a computationally efficient way for producing the clustering tree. Instead, we use the splitting procedure to understand the large sample behavior of the sample clustering procedure. It is reasonable to expect that, as n tends to infinity, the sequence of splits in the sample procedure should resemble the sequence of splits in an analogous procedure defined on the

population. The population procedure can be defined by replacing the averages with the corresponding conditional means. The formal definition is given in the next section.

2.2 Population Clustering Procedure

For the remainder of the paper we assume that the underlying distribution has a real valued density, f. For concreteness, we focus on the case where the support of the distribution is of the form (L_0, R_0), where $-\infty \leq L_0 < R_0 \leq \infty$. Thus, every open interval in (L_0, R_0) contains positive probability. Given an interval $(l, r) \subseteq (L_0, R_0)$, we write $\mu_{l,r}$ for the population conditional mean on (l, r),

$$\mu_{l,r} = \left(\int_l^r f(x)dx \right)^{-1} \int_l^r x f(x)dx. \tag{4}$$

We set $\mu_{r,r} = r$, by continuity. Given an interval $(L, R) \subseteq (L_0, R_0)$, we define

$$G_{L,R}(a) = \mu_{a,R} - \mu_{L,a}, \tag{5}$$

for $a \in [L, R]$. Note that $G_{L,R}(L) = \mu_{L,R} - L$ and $G_{L,R}(R) = R - \mu_{L,R}$.

According to the results in Section 2.1, the sample clustering procedure determines the split partition of a cluster by maximizing the distance between the empirical sub-cluster means. We define the *population clustering procedure* by analogy. Given a cluster (L, R), the population procedure chooses the split that maximizes the distance between the population sub-cluster means. More formally, it finds a point s that maximizes $G_{L,R}$, then partitions (L, R) into subintervals (L, s) and (s, R), on which the procedure is repeated. If s is an interior point of (L, R), we call it a *split point*, and we call the corresponding partition a *split*. Otherwise, the population procedure essentially wants to split off an endpoint, which forces the cluster to be truncated rather than split (see Figure 1 for an illustration). A formal definition of the corresponding *truncation* operation is given in the next subsection.

2.3 Properties of the Population Procedure

In this section we discuss existence and uniqueness of the population clustering procedure. The proofs of the corresponding results are provided in the Supplementary Material. We

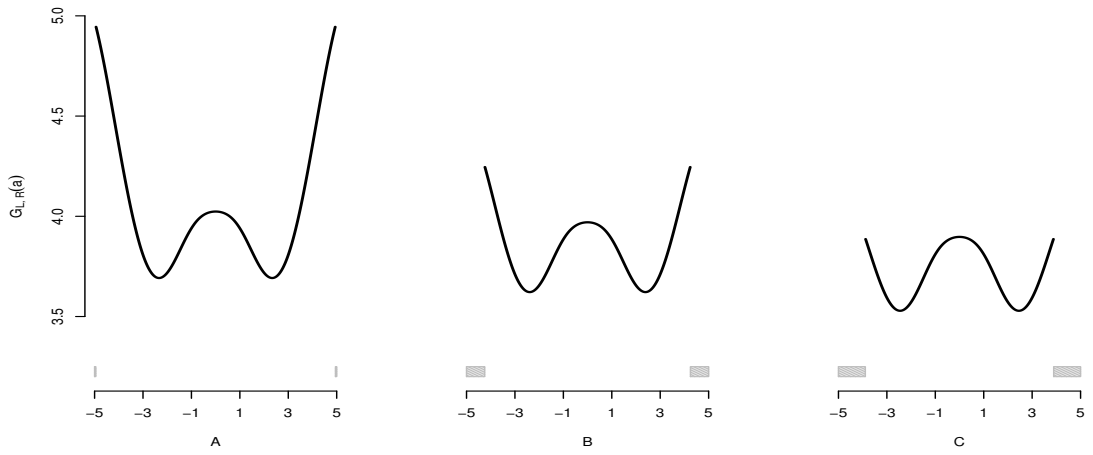

Figure 1: Population criterion function, $G_{L,R}$, corresponding to the Gaussian mixture $0.5\,N(-2,1) + 0.5\,N(2,1)$, is plotted for three choices of (L,R), such that $R = -L$ and $G_{L,R}(L) = G_{L,R}(R) = \max G_{L,R}$. The population clustering procedure continuously truncates the support of the distribution, symmetrically on both sides, until $\max G_{L,R}$ can be achieved at an interior point (plot C). Then, the procedure places a split point at zero.

start with a formal definition of the *truncation* operation in the population procedure. Given a cluster (L,R),

(i) if $\arg\max G_{l,R} = \{l\}$ for all $l \in [L, L^*)$, and $\arg\max G_{L^*,R} \neq \{L^*\}$, then the population procedure truncates (L,R) down to the interval (L^*, R);

(ii) if $\arg\max G_{L,r} = \{r\}$ for all $r \in (R^*, R]$, and $\arg\max G_{L,R^*} \neq \{R^*\}$, then (L,R) is truncated down to (L, R^*);

(iii) if there exists a continuous decreasing function $l \mapsto R_l$, such that $\arg\max G_{l,R_l} = \{l, R_l\}$ for all $l \in [L, L^*)$, and $\arg\max G_{L^*,R_{L^*}} \neq \{L^*, R_{L^*}\}$, then (L,R) is truncated down to (L^*, R_{L^*}).

We first consider an important special case, where the underlying distribution is unimodal. More specifically, suppose the density f is either strictly monotone on its support, (L_0, R_0), or there exists a point c for which f is strictly increasing on (L_0, c) and strictly decreasing on (c, R_0). The following result shows that in this setting, under an additional continuity assumption on f, the population procedure is unique and does not reveal any clusters.

Proposition 3. *If f is continuous and unimodal, then the population clustering procedure is uniquely defined and produces no splits.*

The following simple regularity condition ensures existence of a population clustering procedure with finitely many splits, as we demonstrate in the proof of Proposition 5 below.

C1. Density f is nonzero and differentiable on (L_0, R_0). It has finitely many modes and, at each of its interior modes, admits a non-constant Taylor approximation.

> **Remark.** The last requirement means the following: for each interior mode c, there exists a positive integer k, such that f is k times differentiable at c with $f^{(k)}(c) \neq 0$.

To address the question of uniqueness, consider the following counterexample. Suppose the underlying distribution is uniform on (L, R). Then, the criterion function $G_{L,R}$ is constant on its domain, and the population procedure applied to cluster (L, R) may place a split point anywhere in its interior. Thus, there are infinitely many versions of the population clustering procedure. The following regularity condition explicitly rules out such settings, by requiring that the interior of (L, R) contains at most one maximizer of $G_{L,R}$.

C2. When the population procedure performs a split, the location of the split point is uniquely determined.

Note that condition C2 holds for each distribution with a continuous unimodal density, as a direct consequence of Proposition 3. The next result shows that C2 also holds for all bimodal densities, provided the smoothness condition, C1, is satisfied.

Proposition 4. *If the underlying distribution is bimodal, then condition C1 implies condition C2.*

The following proposition establishes existence and uniqueness of the population clustering procedure in the general setting. In Section 3 we show that under the same regularity conditions the sample procedure consistently estimates its population counterpart.

Proposition 5. *If regularity conditions C1 and C2 are satisfied, then the population clustering procedure is uniquely defined and produces finitely many splits.*

In Section D of the Supplementary Material we document the behaviour of the population clustering procedure for various mixtures of Gaussian distributions on the real line.

3 Main Results

In this section we show that the clustering tree produced by criterion (2) consistently estimates the clustering tree produced by the corresponding population procedure. We also derive the corresponding rates of convergence, and propose a novel post-processing modification of the sample clustering procedure.

3.1 Consistency

We start with some useful notation. In both the sample and the population, each split is characterized by a triple (L, s, R), where interval (L, R) is the cluster being split, and s is the split point, located inside (L, R). We write $P_{L,R}$ for the probability assigned to the interval (L, R) by the underlying distribution. For a split $sp = (L, s, R)$ we define its size as $size(sp) = \min\{P_{L,s}, P_{s,R}\}$. When the probabilities in the above definition are replaced with the corresponding sample frequencies, we write $\widehat{size}(sp)$ for the resulting quantity, and refer to it as the empirical size.

The set of all the population splits, denoted by $\mathcal{S}$, defines the population clustering tree. Similarly, the set of all sample splits, $\widehat{\mathcal{S}}$, defines the sample tree. The cardinality of $\widehat{\mathcal{S}}$ tends to infinity as the sample size grows. Alternatively, according to Proposition 5, under mild regularity conditions the population procedure produces finitely many splits, together with some truncations. To establish consistency, we divide the sample splits into "big" and "small", based on their empirical size, then show that the first group converges to the population splits, while the second is asymptotically negligible. The formal definition is given below. We write d_H for the Hausdorff distance between subsets of a Euclidean space.

Definition 1. *Write $\widehat{\mathcal{S}}_\alpha$ for $\{sp \in \widehat{\mathcal{S}} : \widehat{size}(sp) > \alpha\}$ and let α^* be the smallest split size in the population procedure. We call the sample clustering procedure* ***strongly consistent*** *if, for each α in $(0, \alpha^*)$, the following statements hold almost surely,*

$$|\widehat{\mathcal{S}}_\alpha| \to |\mathcal{S}| \tag{6}$$

$$d_H(\widehat{\mathcal{S}}_\alpha, \mathcal{S}) \to 0 \qquad \textit{and} \tag{7}$$

$$\max\{size(sp) : sp \in \widehat{\mathcal{S}} \setminus \widehat{\mathcal{S}}_\alpha\} \to 0. \tag{8}$$

If we replace almost sure convergence with convergence in probability, we have a weaker

notion of consistency, which holds automatically when the sample procedure is strongly consistent. In particular, displays (6) and (7) imply that, except on a set of probability tending to zero, there is a one to one correspondence between $\widehat{\mathcal{S}}_\alpha$ and the set of all population splits, such that each split in $\widehat{\mathcal{S}}_\alpha$ converges to its population counterpart with respect to the usual Euclidean distance. The next result, which is proved in the Supplementary Material, establishes consistency of the sample procedure.

Theorem 1. *Suppose that regularity conditions C1 and C2, given in Section 2.3, are satisfied. Then, the sample clustering procedure is strongly consistent.*

> **Remark.** It follows from the proof that the result continues to hold if we replace $\widehat{size}$ with $size$ in the definition of $\widehat{\mathcal{S}}_\alpha$ and/or replace $size$ with $\widehat{size}$ in display (8).

Consider the important case of a continuous unimodal density f. Proposition 3 in Section 2.3 and the proof of Theorem 1 imply that in this case condition C1 is not required for consistency. Note also that the population procedure produces no splits, by Proposition 3. It follows that all of the sample splits are uniformly asymptotically negligible.

Corollary 1. *If f is continuous and unimodal, then the maximum size of all the splits in the sample clustering procedure goes to zero almost surely.*

In the next section we extend the results in Theorem 1 by establishing the rates of convergence for the sample clustering procedure.

3.2 Rates of Convergence

To establish the rates of convergence for the sample splitting procedure, we need an additional regularity condition. We use the term *population cluster* to refer to all intervals that appear along the path of the population procedure.

C3. For each population cluster (L, R) and each $t \in \arg\max_{[L,R]} G_{L,R}$, if $t \in (L, R)$, then $G''(t) \neq 0$, otherwise $G'_{L,R}(t) \neq 0$.

The requirement on $G_{L,R}(t)''$, imposed for each population split (L, t, R), is the standard M-estimation assumption that requires the second derivative of the population criterion function to be nonsingular at the population maximum (e.g. van der Vaart and Wellner 1996; van der Vaart 1998). The requirement on $G'_{L,R}$ is the analog of the aforementioned

M-estimation assumption in the case where the population criterion function, $G_{L,R}$, is maximized at an endpoint of $[L, R]$, rather than in the interior. In this case, the behaviour of $G_{L,R}$ near its maximum is characterized by the first derivative, rather than by the second.

Let $\widehat{\mathcal{S}}(\tau)$ contain all the sample splits $\widehat{sp} = (L, \widehat{t}, R)$, for which the sample frequencies of (L, R), $(-\infty, L)$ and (R, ∞) are greater than or equal to τ. In Theorem 2 we restrict our attention to the sample splits in $\widehat{\mathcal{S}}(\tau)$, for arbitrarily small but positive τ. Without this restriction, the rate of convergence in (10) would change. In particular, split sizes larger than $O_p(\log n/n)$ are produced when the sample procedure is applied to intervals whose widths tend to zero. Also, larger split sizes may appear near the boundary of the support of the distribution. The general approach used in the proof of Theorem 2 would also establish these slower rates of convergence. However, the exact form of the new rates depends on the behavior of the density near the boundary of the support and on the aforementioned intervals of negligible width. Instead, we present a clean result, with just one rate of convergence for all the small sample splits, while only imposing some simple regularity conditions.

Theorem 2. *Suppose that regularity conditions C1-C3 are satisfied. Let α^* be the smallest split size in the population procedure. Then, for each α in $(0, \alpha^*)$,*

$$d_H(\widehat{\mathcal{S}}_\alpha, \mathcal{S}) = O_p\left(n^{-1/3}\right) \qquad \textit{and} \tag{9}$$

$$\max\{size(sp) : sp \in (\widehat{\mathcal{S}} \setminus \widehat{\mathcal{S}}_\alpha) \cap \widehat{\mathcal{S}}(\tau)\} = O_p\left(\log n/n\right). \tag{10}$$

Remark. As we point out in the proof, if the domain, (L_0, R_0), of the underlying distribution is bounded, then we can remove the lower bounds on the sample frequencies of $(-\infty, L)$ and (R, ∞) from the definition of $\widehat{\mathcal{S}}(\tau)$ by assuming, instead, that $f(L_0)$ and $f(R_0)$ are nonzero.

The proof of Theorem 2 is provided in the Supplementary Material. The intuition for the presented rates of convergence can be described, informally, as follows. We bound the distance between the sample split point, $\widehat{t}$, and its population counterpart, t, by characterizing the behaviour of the sample criterion function near t. The sample criterion, $\widehat{G}_{L,R}(a)$, is the empirical analog of the population criterion, $G_{L,R}(a)$, and is defined as the difference between the averages of the observations in $[a, R]$ and $[L, a]$, respectively. We examine the decrease in $G_{L,R}(t)$ that occurs when t is perturbed by a small amount, δ. Then, we contrast this decrease with the stochastic term given by the difference between the corresponding

deviations in $\widehat{G}_{L,R}$ and $G_{L,R}$. The order of this term is roughly $\sqrt{\delta/n}$. When t is a population split point and, thus, lies in the interior of (L, R), the corresponding decrease in $G_{L,R}$ is quadratic in δ. Balancing out the two terms yields the cube root asymptotic behaviour (cf. Kim and Pollard 1990) for the sample split point. When function $G_{L,R}$ is maximized at an endpoint of the interval $[L, R]$, as in the case of truncation, the decrease in $G_{L,R}$ is linear in δ. Balancing this decrease with the stochastic term of order $\sqrt{\delta/n}$ suggests that the sample split point is a $O_p(n^{-1})$ amount away from the boundary of (L, R). Uniformity of the rate over all the small sample splits requires an additional $\log n$ factor.

In the next section we take advantage of our asymptotic results to propose a key modification to the sample clustering procedure.

3.3 Big Merge Tracker: Post-processing the Sample Procedure

Theorem 1 demonstrates that the sample analog of the truncation operation is peeling a large number of tiny clusters off the ends of a large cluster. It follows that when recording sample splits we should distinguish between those that correspond to splits in the population procedure and those that correspond to truncations. Based on this observation, we propose to post-process the sample clustering procedure by only keeping the splits with significant empirical sizes. More specifically, given a threshold α, if the cardinality of one of the sub-clusters is below αn, the corresponding split is removed from the final output.

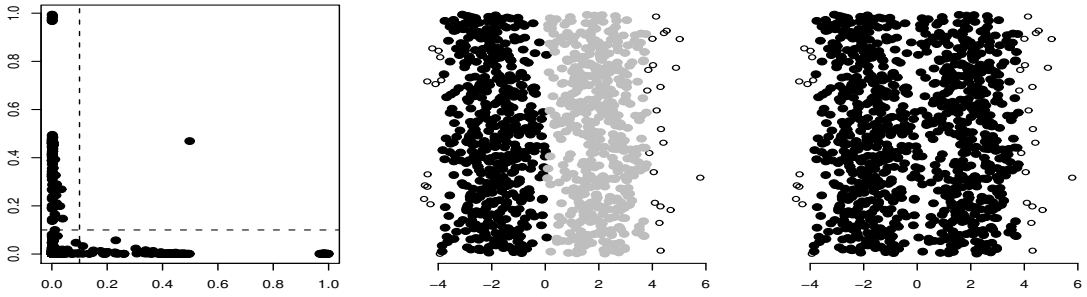

Figure 2: The plots illustrate the path of Algorithm 1 on a sample of 1000 observations from a symmetric Gaussian mixture, $0.5\,N(-2, 1) + 0.5\,N(2, 1)$. The scatter plot on the left displays the sample frequencies for each pair of clusters merged along the path. The next two plots show the cluster memberships before and after the *big merge*.

Figure 2 illustrates the path of Algorithm 1 on a sample of 1000 observations, generated independently from the symmetric Gaussian mixture distribution used in Figure 1. The

scatter plot on the left displays the sample frequencies for each pair of clusters merged along the path. We found only one merge in which both clusters pass the $\alpha = 0.1$ threshold. The *big merge* occurs at a point where the current number of clusters is 32. The two rightmost plots display cluster memberships before and after the merge. The non-shaded points belong to clusters with non-appreciable size.

The equivalence between the splitting procedure and the merging algorithm implies that in the post-processing step of Algorithm 1 we only keep the merges where the cardinality of each of the merging clusters is above αn. For any such merge, we place the split point midway between the two closest representatives of the two clusters being merged. We also replace the stored merges with the corresponding split points. The resulting sequence of split points can then be reinterpreted as a sequence of splits, or a sequence of merges, using the full sample. For example, if the final output contains no split points, then all of the observations in the sample are placed in the same cluster. We call this modified approach the *Big Merge Tracker* (BMT) with threshold α. It is a direct consequence of Theorems 1 and 2 that, under the respective regularity conditions, the BMT split points are strongly consistent and converge at the $O_p(n^{-1/3})$ rate. In the next section we analyze the empirical performance of the Big Merge Tracker.

4 Simulation Study and Real Data Analysis

In this section we show strong performance of the proposed BMT approach relative to popular existing methods for assessing modality and estimating the number of clusters. We also illustrate the use of our methodology in analysis of single-cell virology datasets. In addition, in Section E.3 of the Supplementary Material we apply BMT on very large simulated data sets and demonstrate its superior scalability properties.

When the separation between two clusters is very small, the population splitting procedure can still be successful at finding a split point by massively truncating the support. This *zooming-in* effect may result in larger sizes of the *small* sample splits, as we discussed, from a theoretical standpoint, in the paragraph above the statement of Theorem 2. To counteract this phenomenon, we propose an *adjustment* to the Big Merge Tracker. If the sum of the sample frequencies for the two merging clusters in the last big merge is less than 50%, we do not report any merges. Preventing the corresponding splitting procedure from

truncating more than 50% of the data, while searching for the first split, slightly reduces its efficiency, but makes it more robust to sampling fluctuations. Throughout this section we use the adjustment described above and set the BMT threshold, α, equal to 0.1 (Algorithm 3 in Section E of the Supplementary Material contains the full pseudocode).

4.1 Modality Assessment

Testing for homogeneity of a population is an important statistical problem (Aitkin and Rubin, 1985; Müller and Sawitzki, 1991; Roeder, 1994). Here, we use the BMT to detect the presence of two or more dominant modes in the density. In Table 1 we compare our approach with two popular modality assessment procedures: (i) kernel density estimate based test of Silverman (1981) (ii) histogram based Diptest proposed by Hartigan and Hartigan (1985). P-values of the Silverman test are calculated using the R-package referenced in Vollmer *et al.* (2013). R-packge of Maechler (2013) is used for implementing the Dip test. Detailed descriptions of these procedures are given in Section E.1 of the Supplementary Material.

We consider 5 different simulation scenarios, in which 100 independent samples of size 10000 were generated and subjected to modality analysis. Table 1 reports the percentage of cases in which multi-modality was detected. P-values for the Dip and Silverman tests were computed based on 1000 MCMC simulations, and decision on the null hypothesis of unimodality was made at the 5% level of significance. The mean and the standard deviation of the p-values from these tests are also reported. In the two unimodal scenarios the BMT is on par with the Silverman and the Dip tests in confirming unimodality of the population distribution with high certainty. In the three non-unimodal cases, which include normal and beta mixtures, the BMT shows better performance in detecting multi-modality.

4.2 Estimating the Number of Clusters

We study the potency of the BMT in detecting the true number of clusters. We compare its performance with the following *number of clusters* estimation methods: (i) the CH index of Caliński and Harabasz (1974) (ii) the KL index of Krzanowski and Lai (1988) (iii) the H measure of Hartigan (1975) (iv) the Silhouette statistic based KR index of Kaufman and Rousseeuw (2009), (v) the Gap statistic of Tibshirani *et al.* (2001) (vi) the Jump

Population Density	Dip Test P-value (D)			Silverman Test P-value (S)			BMT
	Mean (D)	Std(D)	**% multi-mode**	Mean (S)	Std(S)	**% multi-mode**	**% multi-mode**
N(0,1)	0.99	0.04	0.00	0.48	0.25	0.00	0.00
Beta(2,4)	0.98	0.04	0.00	0.54	0.28	2.00	0.00
$\{N(-1.1,1)+N(1.1,1)\}/2$	0.81	0.22	0.00	0.22	0.21	29.00	69.00
$\{Beta(4,6)+Beta(7,3)\}/2$	0.84	0.22	0.00	0.31	0.25	21.00	49.00
$\{N(-2.5,1)+N(0,1)+N(2.5,1)\}/3$	0.10	0.14	52.00	0.03	0.03	79.00	96.00

Table 1: Simulation study to compare multi-modality detection methods

statistic of Sugar and James (2003) (vii) the clustering prediction strength criterion of Tibshirani and Walther (2005), and (viii) the bootstrap based cluster instability minimizing criterion of Fang and Wang (2012), which is inspired by the stable clusters selection approach of Wang (2010). Detailed descriptions of these procedures are provided in Section E.2 of the Supplementary Material. We consider one multivariate and five univariate regimes. 100 independent replications with the sample size of 5000 are used in each simulation setting, and the distribution of the number of clusters detected by each method is reported. The eight competitor methods are implemented using a number of different clustering approaches via the `NbClust` R-package of Charrad *et al.* (2014) and the `fpc` package of Hennig (2014). More specifically, we use k-means clustering with the Euclidean distance metric (the corresponding results are reported in Table 2), as well as Ward's method (Ward Jr, 1963), Centroid-based clustering (Kaufman and Rousseeuw, 2009), PAM (Kaufman and Rousseeuw, 1990), CLARA (Kaufman and Rousseeuw, 1986), clustering by merging Gaussian mixture components (Hennig, 2010) and hierarchical clustering initialized Gaussian mixture model based clustering method of Fraley and Raftery (2002). All the results for the non-k-means clustering approaches are reported in the tables provided in Section E.2 of the Supplementary Material.

In our first univariate example, we consider a non-symmetric mixture of two normal densities. Each of them has unit variance and their means are fairly well-separated. We observe that the CH, KL and Hartigan methods struggle to recover the bimodal structure, while the others successfully detect the two clusters (we note that in this setting, the CH index performs better when it uses the centroid based clustering algorithm). The next three simulation scenarios correspond to non-symmetric tri-modal population densities that

True Population Density	Methods	Number of Clusters									
		1	2	3	4	5	6	7	8	9	10+
$0.3\,N(-4,1)+0.7\,N(4,1)$	CH	0	**29**	12	4	2	2	4	11	14	22
	KL	0	**39**	10	5	5	8	8	8	9	8
	Hartigan	0	**0**	32	16	12	10	11	5	7	7
	Silhouette	0	**100**	0	0	0	0	0	0	0	0
	Gap	0	**100**	0	0	0	0	0	0	0	0
	Jump	0	**99**	0	0	0	0	0	0	0	1
	Pred Str.	0	**100**	0	0	0	0	0	0	0	0
	Stability	0	**100**	0	0	0	0	0	0	0	0
	BMT	0	**93**	7	0	0	0	0	0	0	0
$0.3\,N(-3,1)+0.35\,N(0,1)+0.35\,N(3,1)$	CH	0	0	**0**	0	1	0	6	13	28	52
	KL	0	11	**8**	11	9	6	14	12	19	10
	Hartigan	0	0	**69**	13	7	4	2	4	1	0
	Silhouette	0	8	**92**	0	0	0	0	0	0	0
	Gap	0	100	**0**	0	0	0	0	0	0	0
	Jump	0	0	**100**	0	0	0	0	0	0	0
	Pred Str.	0	0	**100**	0	0	0	0	0	0	0
	Stability	0	11	**89**	0	0	0	0	0	0	0
	BMT	0	5	**95**	0	0	0	0	0	0	0
$0.3\,t_1(-3)+0.35\,t_1(0)+0.35\,t_1(3)$	CH	0	7	**9**	7	11	7	6	10	14	29
	KL	0	18	**23**	15	14	5	4	9	6	6
	Hartigan	0	0	**46**	24	11	5	4	5	1	4
	Silhouette	0	71	**26**	3	0	0	0	0	0	0
	Gap	0	36	**64**	0	0	0	0	0	0	0
	Jump	17	11	**9**	13	12	5	11	9	7	6
	Pred Str.	0	24	**53**	19	4	0	0	0	0	0
	Stability	0	69	**30**	0	1	0	0	0	0	0
	BMT	0	1	**99**	0	0	0	0	0	0	0
$0.3\,\text{dexp}(-3)+0.35\,\text{dexp}(0)+0.35\,\text{dexp}(3)$	CH	0	0	**0**	0	1	0	6	14	23	56
	KL	0	13	**8**	18	9	11	7	20	10	4
	Hartigan	0	0	**55**	19	9	10	3	1	1	2
	Silhouette	0	39	**61**	0	0	0	0	0	0	0
	Gap	0	100	**0**	0	0	0	0	0	0	0
	Jump	0	0	**100**	0	0	0	0	0	0	0
	Pred Str.	0	0	**100**	0	0	0	0	0	0	0
	Stability	0	0	**100**	0	0	0	0	0	0	0
	BMT	0	0	**100**	0	0	0	0	0	0	0
$\{Beta(8,2)+Beta(5,5)+Beta(2,8)\}/3$	CH	0	0	**2**	0	3	3	6	16	22	48
	KL	0	13	**3**	11	9	10	15	16	9	14
	Hartigan	0	0	**57**	14	9	6	5	5	0	4
	Silhouette	0	0	**100**	0	0	0	0	0	0	0
	Gap	0	100	**0**	0	0	0	0	0	0	0
	Jump	99	0	**1**	0	0	0	0	0	0	0
	Pred Str.	0	78	**17**	4	1	0	0	0	0	0
	Stability	0	0	**25**	60	15	0	0	0	0	0
	BMT	0	0	**100**	0	0	0	0	0	0	0
$\{0.5\,N(-2,1)+0.5\,N(2,1)\}$ $\otimes\,N(0,1)$ $\otimes\,N(0,1)$ $\otimes\,\chi^2_1$ $\otimes\,\chi^2_1$	CH	0	**0**	1	0	4	5	11	10	23	46
	KL	0	**18**	11	14	7	13	9	11	11	6
	Hartigan	0	**0**	52	16	13	4	7	4	4	0
	Silhouette	0	**0**	87	12	0	1	0	0	0	0
	Gap	0	**100**	0	0	0	0	0	0	0	0
	Jump	0	**14**	0	62	0	0	0	0	8	16
	Pred Str.	0	**86**	1	13	0	0	0	0	0	0
	Stability	0	**85**	0	15	0	0	0	0	0	0
	BMT	0	**96**	4	0	0	0	0	0	0	0

Table 2: Number of clusters detected in 100 trials for six simulation scenarios

are mixtures of standard normals, non-central t-densities with one degree of freedom and double exponential densities with the unit rate parameter, respectively. The medians of the mixing densities and the mixture weights match across the three settings, and the separation between the adjacent medians is not large. In these three simulation settings, the Gap statistic approach, as well as the CH, KL and Hartigan measures, has difficulty detecting the true number of clusters. The Silhouette method performs well for Gaussian mixtures but has difficulties in the other two cases. Jump statistic, prediction strength and bootstrap stability approaches do well in the normal and the double exponential cases, however, they do not show good performance in the considerably thicker-tailed case of the mixture of t-densities. For our fifth simulation setting we consider a bounded population density that is a mixture of three Beta densities. Here, only the BMT and the Silhouette do well in recovering the true number of clusters. In our last example we consider a 5 dimensional data set, which is generated from a product density. The first dimension is generated from a symmetric mixture of two Gaussians; the next two dimensions contain white noise, while the forth and fifth dimensions are generated from a central chi-square distribution with one degree of freedom. We observe that, together with the CH, KL and Hartigan measures, the Silhouette and the Jump approaches do not perform well in detecting the two clusters in this data.

BMT does consistently well across each of the six simulation scenarios, outperforming all the other approaches overall. The prediction strength and the cluster stability methods, which are modern state of the art approaches, deliver the best results among the competitors. However, these two methods have significant trouble in the cases of the beta and the non-central t mixtures. When implemented with various non-k-means clustering approaches (see Section E.2 of the Supplementary Material), neither of the competitors considerably improves the performance reported in Table 2. Figure S3 in Section E.2 of the Supplementary Material provides plots of the densities used in the numerical experiments.

4.3 Sub-population Analysis in Single Cell Virology

We demonstrate an application of our clustering approach in a *single-cell Mass Cytometry* (Bendall *et al.*, 2011) based virology study. We analyze the data reported in Sen *et al.* (2014), where the effect of Varicella Zoster Virus (VZV) on human tonsil T cell is studied. VZV is a human herpesvirus and causes varicella and zoster (Zerboni *et al.*, 2014).

We study protein expressions from five independent experiments, each containing an Uninfected (UN) and a Bystander (BY) populations. Bystanders are cells in the VZV infected population, which are not directly infected by the virus, but are influenced by neighboring virus infected cells. Protein expression values are studied on the arcsinh scale. Non-expressed values are uniformly distributed between $[-1, 1]$. Cellular sub-populations are detected by clustering the populations based on the expressions of "core-proteins", which are associated with T cell activation (Newell *et al.*, 2012). Most of the samples have large sizes, usually on the order of $\sim 10^5$. Traditional clustering techniques fail to accommodate such large sample sizes and resort to sub-sampling based approaches (Qiu *et al.*, 2011; Linderman *et al.*, 2012). The BMT, on the other hand, has the advantage of being scalable enough to conduct clustering analysis on the entire sample.

	Experiment I		Experiment II		Experiment III		Experiment IV		Experiment V	
SUB-POPULATIONS	UN	BY	UN	BY	UN	BY	UN	BY	UN	BY
DUAL POSITIVE	8411 (8.8%)	6596 (7.3%)	5253 (5.8%)	4169 5.7%)	4971 (6.0%)	2703 (5.4%)	3795 (5.0%)	1510 (4.6%)	8047 (8.5%)	5225 (8.0%)
DUAL NEGATIVE	2723 (2.8%)	2973 (3.3%)	3537 (3.9%)	2631 (3.6%)	4433 (5.3%)	2935 (5.9%)	4354 (5.8%)	2196 (6.7%)	5012 (5.3%)	2881 (4.4%)
CD4 NON-NAIVE	7993 (8.4%)	10636 (11.8%)	15144 (16.7%)	11556 (15.9%)	21444 (25.9%)	12429 (25.0%)	22149 (29.7%)	8508 (26.1%)	30034 (31.9%)	20469 (31.3%)
CD4 NAIVE	69977 (73.7%)	64119 (71.1%)	57744 (63.7%)	47374 (65.1%)	45764 (55.3%)	27987 (56.3%)	35458 (47.5%)	16390 (50.4%)	40398 (43.0%)	28524 (43.7%)
CD8 NAIVE	5654 (6.0%)	5671 (6.3%)	8599 (9.5%)	6571 (9.0%)	5798 (7.0%)	3490 (7.0%)	8271 (11.1%)	3774 (11.6%)	9869 (10.5%)	7829 (12.0%)
POPULATION SIZE	94837	90157	90641	72699	82637	49672	74540	32497	93878	65244

Table 3: Sizes and Proportions of dominant clusters detected by BMT across 5 independent Virology experiments

We treat three proteins, CD4, CD8 and CD45RA (naive), as core-proteins, as they are typically used to classify T cells. For each of the 10 samples (UN and BY from experiments I-V), based on the expressions of the above three proteins, we performed automated clustering by using BMT in the three dimensional space. Figure S4 and S5 in Section E.4 of the Supplementary Material show that in all the cases, the BMT detects unimodality for CD4 and CD45RA and bimodality for CD8 expression values. Using the bi-modality of CD8 and the BMT detected splits, we classify cells as CD8-high and CD8-low. Also, considering the expression and non-expressions of the other two markers we simultane-

ously classify cells into the following clusters, or sub-populations: (i) Dual positive: CD4 expressed and CD8 High (ii) Dual-negative: CD4 non-expressed and CD8 low (iii) CD4 Non-Naive: CD4 expressed and CD45RA non-expressed (iv) CD4 Naive: CD4 expressed and CD45RA non-expressed (v) CD8 Naive: CD8 high and CD45RA expressed.

Table 3 reports the sizes and proportional representations of these sub-populations, across the five experiments. BMT sub-populations resemble the T-cell biology based phenotypic classification in Sen *et al.* (2014). They also revalidate that the sub-population distribution in the Bystander cells is not much different from that of the uninfected, though the UN sub-population distribution varies across experiments. Using this BMT based categorization of the T cells, sub-population level cell-signaling patterns can be subsequently studied. Figure S6 in Section E.4 of the Supplementary Material shows the heatmap of the protein expressions (core + signaling proteins) of the sub-populations from Experiment I.

5 Acknowledgement

Radchenko's research was partially supported by NSF Grant DMS-1209057. Computing resources were provided by the University of Southern California Center for High Performance Computing. The R code and the data sets used for producing the numerical results of Section 4 can be downloaded from `http://www-bcf.usc.edu/~gourab/code-bmt`.

References

Aitkin, M. and Rubin, D. B. (1985). Estimation and hypothesis testing in finite mixture models. *Journal of the Royal Statistical Society. Series B (Methodological)* 67–75.

Bach, F. R. and Harchaoui, Z. (2008). Diffrac: a discriminative and flexible framework for clustering. In *Advances in Neural Information Processing Systems*, 49–56.

Bendall, S. C., Simonds, E. F., Qiu, P., Amir, E., Krutzik, P. O., Finck, R., Bruggner, R. V., Melamed, R., Trejo, A., Ornatsky, O. I., Balderas, R. S., Plevritis, S. K., Sachs, K., Pe'er, D., Tanner, S. D., and Nolan, G. P. (2011). Single-cell mass cytometry of differential immune and drug responses across a human hematopoietic continuum. *Science (New York, N.Y.)* **332**, 6030, 687–696.

Bondell, H. D. and Reich, B. J. (2008). Simultaneous regression shrinkage, variable selection, and supervised clustering of predictors with oscar. *Biometrics* **64**, 1, 115–123.

Caliński, T. and Harabasz, J. (1974). A dendrite method for cluster analysis. *Communications in Statistics-theory and Methods* **3**, 1, 1–27.

Charrad, M., Ghazzali, N., Boiteau, V., and Niknafs, A. (2014). NbClust: An R package for determining the relevant number of clusters in a data set. *Journal of Statistical Software* **61**, 6, 1–36.

Chi, E. C. and Lange, K. (2013). Splitting methods for convex clustering. *arXiv preprint arXiv:1304.0499* .

Fang, Y. and Wang, J. (2012). Selection of the number of clusters via the bootstrap method. *Computational Statistics & Data Analysis* **56**, 3, 468–477.

Fraley, C. and Raftery, A. E. (2002). Model-based clustering, discriminant analysis, and density estimation. *Journal of the American statistical Association* **97**, 458, 611–631.

Hartigan, J. (1978). Asymptotic distributions for clustering criteria. *The Annals of Statistics* 117–131.

Hartigan, J. A. (1975). *Clustering algorithms*. Wiley.

Hartigan, J. A. and Hartigan, P. (1985). The dip test of unimodality. *The Annals of Statistics* 70–84.

Hennig, C. (2010). Methods for merging gaussian mixture components. *Advances in data analysis and classification* **4**, 1, 3–34.

Hennig, C. (2014). *fpc: Flexible procedures for clustering*. R package version 2.1-9.

Hocking, T., Vert, J.-P., Bach, F., and Joulin, A. (2011). Clusterpath: an algorithm for clustering using convex fusion penalties. In L. Getoor and T. Scheffer, eds., *ICML*, 745–752. Omnipress.

Hoefling, H. (2010). A path algorithm for the fused lasso signal approximator. *Journal of Computational and Graphical Statistics* **19**, 4, 984–1006.

Jain, A. K. (2010). Data clustering: 50 years beyond k-means. *Pattern Recognition Letters* **31**, 8, 651–666.

Kaufman, L. and Rousseeuw, P. (1986). *Clustering large data sets*. Elsevier.

Kaufman, L. and Rousseeuw, P. J. (1990). Partitioning around medoids (program pam). *Finding groups in data: an introduction to cluster analysis* 68–125.

Kaufman, L. and Rousseeuw, P. J. (2009). *Finding groups in data: an introduction to cluster analysis*, vol. 344. John Wiley & Sons.

Ke, T., Fan, J., and Wu, Y. (2013). Homogeneity in regression. *arXiv preprint arXiv:1303.7409* .

Kim, J. and Pollard, D. (1990). Cube root asymptotics. *The Annals of Statistics* 191–219.

Krzanowski, W. J. and Lai, Y. (1988). A criterion for determining the number of groups in a data set using sum-of-squares clustering. *Biometrics* 23–34.

Li, J. (2005). Clustering based on a multilayer mixture model. *Journal of Computational and Graphical Statistics* **14**, 3, 547–568.

Li, J., Ray, S., and Lindsay, B. G. (2007). A nonparametric statistical approach to clustering via mode identification. *Journal of Machine Learning Research* **8**, 8, 1687–1723.

Linderman, M. D., Bjornson, Z., Simonds, E. F., Qiu, P., Bruggner, R. V., Sheode, K., Meng, T. H., Plevritis, S. K., and Nolan, G. P. (2012). Cytospade: high-performance analysis and visualization of high-dimensional cytometry data. *Bioinformatics* **28**, 18, 2400–2401.

Lindsten, F., Ohlsson, H., and Ljung, L. (2011). Clustering using sum-of-norms regularization: With application to particle filter output computation. In *Statistical Signal Processing Workshop (SSP), 2011 IEEE*, 201–204. IEEE.

MacQueen, J. *et al.* (1967). Some methods for classification and analysis of multivariate observations. In *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, vol. 1, 281–297. California, USA.

Maechler, M. (2013). *diptest: Hartigan's dip test statistic for unimodality - corrected code*. R package version 0.75-5.

Müller, D. W. and Sawitzki, G. (1991). Excess mass estimates and tests for multimodality. *Journal of the American Statistical Association* **86**, 415, 738–746.

Newell, E. W., Sigal, N., Bendall, S. C., Nolan, G. P., and Davis, M. M. (2012). Cytometry by time-of-flight shows combinatorial cytokine expression and virus-specific cell niches within a continuum of cd8+ t cell phenotypes. *Immunity* **36**, 1, 142 – 152.

Pelckmans, K., De Brabanter, J., Suykens, J., and De Moor, B. (2005). Convex clustering shrinkage. In *PASCAL Workshop on Statistics and Optimization of Clustering Workshop*.

Pollard, D. (1981). Strong consistency of k-means clustering. *The Annals of Statistics* **9**, 1, 135–140.

Pollard, D. (1982). A central limit theorem for k-means clustering. *The Annals of Probability* 919–926.

Qiu, P., Simonds, E. F., Bendall, S. C., Jr., K. D. G., Bruggner, R. V., Linderman, M. D., Sachs, K., Nolan, G. P., and Plevritis, S. K. (2011). Extracting a cellular hierarchy from high-dimensional cytometry data with spade. *Nature Biotechnology* .

Rinaldo, A. *et al.* (2009). Properties and refinements of the fused lasso. *The Annals of Statistics* **37**, 5B, 2922–2952.

Roeder, K. (1994). A graphical technique for determining the number of components in a mixture of normals. *Journal of the American Statistical Association* **89**, 426, 487–495.

Sen, N., Mukherjee, G., Sen, A., Bendall, S. C., Sung, P., Nolan, G. P., and Arvin, A. M. (2014). Single-cell mass cytometry analysis of human tonsil t cell remodeling by varicella zoster virus. *Cell Reports* **8**, 2, 633 – 645.

Shen, X. and Huang, H.-C. (2010). Grouping pursuit through a regularization solution surface. *Journal of the American Statistical Association* **105**, 490.

Shen, X., Huang, H.-C., and Pan, W. (2012). Simultaneous supervised clustering and feature selection over a graph. *Biometrika* **99**, 4, 899–914.

Silverman, B. W. (1981). Using kernel density estimates to investigate multimodality. *Journal of the Royal Statistical Society. Series B (Methodological)* 97–99.

Soltanolkotabi, M. and Candés, E. J. (2012). A geometric analysis of subspace clustering with outliers. *Ann. Statist.* **40**, 4, 2195–2238.

Sugar, C. A. and James, G. M. (2003). Finding the number of clusters in a dataset. *Journal of the American Statistical Association* **98**, 463.

Tan, K. M. and Witten, D. (2015). Statistical properties of convex clustering. *arXiv preprint arXiv:1503.08340* .

Thorndike, R. L. (1953). Who belongs in the family? *Psychometrika* **18**, 4, 267–276.

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)* 267–288.

Tibshirani, R. (2011). Regression shrinkage and selection via the lasso: a retrospective. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **73**, 3, 273–282.

Tibshirani, R., Saunders, M., Rosset, S., Zhu, J., and Knight, K. (2005). Sparsity and smoothness via the fused lasso. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **67**, 1, 91–108.

Tibshirani, R. and Walther, G. (2005). Cluster validation by prediction strength. *Journal of Computational and Graphical Statistics* **14**, 3, 511–528.

Tibshirani, R., Walther, G., and Hastie, T. (2001). Estimating the number of clusters in a data set via the gap statistic. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **63**, 2, 411–423.

Tibshirani, R. J. (2013). Adaptive piecewise polynomial estimation via trend filtering. *arXiv preprint arXiv:1304.2986* .

van der Vaart, A. W. (1998). *Asymptotic Statistics*. Cambridge University Press.

van der Vaart, A. W. and Wellner, J. A. (1996). *Weak Convergence and Empirical Processes: With Applications to Statistics*. Springer-Verlag.

Vollmer, S., Holzmann, H., and Schwaiger, F. (2013). Peaks vs components. *Review of Development Economics* **17**, 2, 352–364.

Wang, J. (2010). Consistent selection of the number of clusters via crossvalidation. *Biometrika* **97**, 4, 893–904.

Ward Jr, J. H. (1963). Hierarchical grouping to optimize an objective function. *Journal of the American statistical association* **58**, 301, 236–244.

Witten, D. M. and Tibshirani, R. (2010). A framework for feature selection in clustering. *Journal of the American Statistical Association* **105**, 490.

Xu, L., Neufeld, J., Larson, B., and Schuurmans, D. (2004). Maximum margin clustering. In *Advances in neural information processing systems*, 1537–1544.

Zerboni, L., Sen, N., Oliver, S. L., and Arvin, A. M. (2014). Molecular mechanisms of varicella zoster virus pathogenesis. *Nature Reviews Microbiology* .

Zhu, C., Xu, H., Leng, C., and Yan, S. (2014). Convex optimization procedure for clustering: Theoretical revisit. In *Advances in Neural Information Processing Systems*, 1619–1627.

Supplementary Material: Consistent clustering using an ℓ_1 fusion penalty

A Proof of Theorem 1

Preliminaries

Here we provide three lemmas that make important contributions to the proof of the theorem. The first allows us to focus on a compact support, the second established uniform convergence of the sample criterion functions on compact sets, and the third derives useful continuity properties of the population solution. The proofs of the lemmas are provided below, after the main argument for Theorem 1. All of the results require regularity conditions C1 and C2, which are given in Section 2.3 of the main paper. In the proof of the lemmas we verify that these conditions can be removed for densities f that are unimodal and continuous. First, we handle the case where the support of density is unbounded. By analogy with the term population cluster, we use *sample cluster* to refer to intervals that appear along the path of the sample clustering procedure. We follow the traditional approach of using ω to denote a general element of the sample space.

Lemma 1. *For every positive ϵ, δ and δ', and almost all ω, there exist a sample cluster $(\widehat{L}, \widehat{R})$, a population cluster (L, R), and a bounded interval B_ϵ, which depends only on ϵ, such that, for all sufficiently large n,*

(a) $P_{-\infty,\widehat{L}} \leq \epsilon$ *and* $P_{\widehat{R},\infty} \leq \epsilon$

(b) $|\widehat{R} - R| \leq \delta$ *and* $|\widehat{L} - L| \leq \delta'$

(c) $(\widehat{L}, \widehat{R}) \subseteq B_\epsilon$.

Remark. Note that L, R, $\widehat{L}$ and $\widehat{R}$ depend on ω, while B_ϵ does not.

To draw connections between the sample and the population clustering procedures, we define the sample criterion function, $\widehat{G}_{L,R}$, as the empirical analog of the population criterion, $G_{L,R}$. More formally, we write $\widehat{\mu}_{l,r}$ for the the sample average on $[l, r]$ and let

$$\widehat{G}_{L,R}(a) = \widehat{\mu}_{a,R} - \widehat{\mu}_{L,a}, \tag{S1}$$

for $a \in [L, R]$, with the following important caveat. If either of the intervals $[L, a]$ and $[a, R]$ contains no observations, we consider $\widehat{G}_{L,R}(a)$ to be *undefined*. We now summarize the sample clustering procedure using the newly defined criterion function. Given a cluster (L, R), the sample procedure finds a point s that maximizes $\widehat{G}_{L,R}$, then splits (L, R) into sub-clusters (L, s) and (s, R). Note that s is guaranteed to be located strictly between two of the observations.

To understand the behavior of the sample clustering procedure on a bounded support, we establish uniform convergence of the of the sample criterion functions to their population counterparts.

Lemma 2. *Define* $A_M = \{(L, R),\ -M \leq L < R \leq M\}$. *Then, as* $n \to \infty$,

$$\sup_{(L,R)\in A_M} \max_{a\in[L,R]} |\widehat{G}_{L,R}(a) - G_{L,R}(a)| \ \to\ 0, \tag{S2}$$

almost surely, for each positive M.

> **Remark.** In the expression $\max_{a\in[L,R]} \widehat{G}_{L,R}(a)$ the maximum is taken over the closed interval $[\,\min\{\boldsymbol{x}_n \cap [L, R]\}\,,\ \max\{\boldsymbol{x}_n \cap [L, R]\}\,]$. By the classical Glivenko-Cantelli theorem, the endpoints of this interval converge to L and R almost surely, and the convergence is uniform over (L, R) in A_M, for each fixed M.

The next result derives some important properties of the population procedure, which we then relate to the sample procedure with the help of Lemma 2. To state the result, we need to modify R_L, which is a function of L that appears in the definition of the two sided truncation. Suppose that the population procedure truncates interval (L_s, R_s) to interval (L^*, R^*) in a two-sided fashion. We will extend function R_L to be defined in open neighborhoods of L_s and L^*, using the equation $G_{L,R_L}(L) = G_{L,R_L}(R_L)$. The proof of Proposition 5 shows that, provided the neighborhoods are small enough, the extension is uniquely determined if we require that function R_L remains continuous and decreasing. To simplify the notation we write $\max G_{L,R}$ for $\max_{a\in[L,R]} G_{L,R}(a)$.

Lemma 3. *Let* (L_s, R_s) *be a bounded cluster, produced by the population procedure, such that* $\arg\max G_{L_s,R_s} = \{L_s, R_s\}$. *Suppose that the first split of the population procedure applied to* (L_s, R_s) *is given by* (L^*, s^*, R^*). *For every sufficiently small positive* ϵ *there exist positive* ϵ', δ *and* δ', *such that* $\epsilon' \leq \epsilon$, $\delta' \leq \delta \leq \epsilon/2$ *and* $|R_L - R_{L+\delta'}| \leq \delta$ *for all* L

in $[L_s - \delta', L^* + \epsilon']$*, and the following sets of inequalities hold as long as* $|R - R_L| \le 2\delta$*:*

$$\max G_{L,R} \;>\; \max_{[L+\delta', R-\delta]} G_{L,R} \qquad \textit{for} \quad L \in [L_s - \delta', L^* - \epsilon'],$$

$$\max G_{L,R} \;>\; \max_{[L+\delta', R-\delta]\setminus(s^*-\epsilon, s^*+\epsilon)} G_{L,R} \qquad \textit{for} \quad L \in [L^* - \epsilon', L^* + \epsilon'/2],$$

$$\max G_{L,R} \;>\; \max_{[L,R]\setminus(s^*-\epsilon, s^*+\epsilon)} G_{L,R} \qquad \textit{for} \quad L \in [L^* + \epsilon'/2, L^* + \epsilon'].$$

$$\max_{[R-\delta, R]} G_{L,R} \;>\; \max_{[L, L+\delta']} G_{L,R} \quad \textit{for} \quad L \in [L_s - \delta', L^* + \epsilon'], \;\; R \in [R_L + \delta, R_L + 2\delta],$$

$$\max_{[L, L+\delta']} G_{L,R} \;>\; \max_{[R-\delta, R]} G_{L,R} \quad \textit{for} \quad L \in [L_s - \delta', L^* + \epsilon'], \;\; R \in [R_L - 2\delta, R_L - \delta].$$

Remark. A minor modification is required when $L_s = L_0$ or $R_s = R_0$, where (L_0, R_0) is the support of f: each interval appearing in the statement needs to be replaced by its intersection with $[L_0, R_0]$.

We will now use Lemma 2 to show that if we replace G with $\widehat{G}$ in the five sets of inequalities above, then for each ω in a set of probability one the resulting sample inequalities simultaneously hold for all sufficiently large n. For example, the first set becomes $\max \widehat{G}_{L,R} > \max_{[L+\delta, R-\delta]} \widehat{G}_{L,R}$. We also allow the cluster (L_s, R_s) to depend on ω, under an additional assumption that there exists a finite deterministic M, for which bound $\max\{|L_s - \delta'|, |R_s + \delta|\} \le M$ holds almost surely. To establish the sample inequalities, we define $\Delta_n = \sup_{A_M} \max |\widehat{G}_{L,R} - G_{L,R}|$ and

$$\tau \;=\; \min_{L \in [L_s - \delta', L^* - \epsilon'], \, |R - R_L| \le 2\delta} \left(\max G_{L,R} - \max_{[L+\delta', R-\delta]} G_{L,R} \right).$$

Continuity of $G_{L,R}(a)$ with respect to a, L an R, together with the first set of inequalities in Lemma 3, implies $\tau > 0$. Taking into account Lemma 2, as well as the continuity of $G_{L,R}(a)$, we derive the following inequalities, which hold for all ω in a set of probability one, and all sufficiently large n:

$$\max \widehat{G}_{L,R} - \max_{[L+\delta, R-\delta]} \widehat{G}_{L,R} > \max G_{L,R} - \max_{[L+\delta, R-\delta]} G_{L,R} - 2\Delta_n - o(1) \ge \tau - o(1) > 0.$$

The argument for the rest of the inequalities is analogous.

Main Body of the Proof

We restrict our attention to the set of probability one that remains after we cast out the negligible sets during the application of the almost sure results from Lemma 1 and Lemma 2. We focus on a fixed ω, but suppress the dependence on ω to simplify the notation.

Consider an arbitrary positive ϵ. We will investigate the behavior of the sample procedure, starting with an appropriately bounded sample cluster, $(\widehat{L}_s, \widehat{R}_s)$, that is close to some population cluster, (L_s, R_s). More specifically, we assume $(\widehat{L}_s, \widehat{R}_s) \subseteq B_\epsilon$, where B_ϵ is a bounded interval that does not depend on ω. We also assume $|\widehat{R}_s - R_s| \leq \delta$ and $|\widehat{L}_s - L_s| \leq \delta'$, for all sufficiently large n, where δ and δ' are to be specified later. We will track the sample and population procedures up to the first population split and show, under only the above assumptions, that the two procedures are close in an appropriate sense. We need to consider the following options for the population procedure applied to (L_s, R_s). First: two-sided truncation, then a split; second: one-sided truncation, then a two sided truncation, followed by a split; third: one-sided truncation, followed by a split. There are three more possibilities, where in each of the initial options the procedure does not produce a split. Finally, the cluster (L_s, R_s) may be split right away.

We first focus on the important case of a two-sided truncation, followed by a split. The rest of the cases can be handled with only minor modifications to the argument. Let $sp = (L^*, s^*, R^*)$ be the first split of the population procedure applied to (L_s, R_s). The population procedure performs a two-sided truncation along (L, R_L), where R_L is a continuous decreasing function of L, such that $R_{L_s} = R_s$ and $R_{L^*} = R^*$. Given our values of ϵ, L_s, R_s and sp, we let ϵ', δ and δ' be the quantities from Lemma 3. Recall that $\epsilon' \leq \epsilon$ and $\delta' \leq \delta \leq \epsilon/2$. As we demonstrate in the paragraph below the statement of Lemma 3, the five sets of population inequalities given in the lemma have natural sample counterparts. The first three sets of inequalities hold for all R with $|R - R_L| \leq 2\delta$ and all sufficiently large n:

$$\begin{aligned}
\max \widehat{G}_{L,R} &> \max_{[L+\delta', R-\delta]} \widehat{G}_{L,R} && \text{for} \quad L \in [L_s - \delta', L^* - \epsilon'] \\
\max \widehat{G}_{L,R} &> \max_{[L+\delta', R-\delta] \setminus (s^*-\epsilon, s^*+\epsilon)} \widehat{G}_{L,R} && \text{for} \quad L \in [L^* - \epsilon', L^* + \epsilon'/2] \\
\max \widehat{G}_{L,R} &> \max_{[L, R] \setminus (s^*-\epsilon, s^*+\epsilon)} \widehat{G}_{L,R} && \text{for} \quad L \in [L^* + \epsilon'/2, L^* + \epsilon'].
\end{aligned} \tag{S3}$$

To simplify the notation, we write $(\widehat{L}, \widehat{R})$ for the current cluster in the sample procedure.

We call a sample split *small* if the split point is placed in $[\widehat{L}, \widehat{L}+\delta']$ or $[\widehat{R}-\delta, \widehat{R}]$ and *big* if the split point is placed in $[s^*-\epsilon, s^*+\epsilon]$. Inequalities (S3) directly imply that the following properties hold for all sufficiently large n, provided $|\widehat{R} - R_{\widehat{L}}| \leq 2\delta$ is satisfied:

(i) if $\widehat{L} \leq L^* - \epsilon'$, then the next split of the sample procedure is small;

(ii) if $\widehat{L} \in [L^* - \epsilon', L^* + \epsilon'/2]$, then the next sample split is either small or big;

(iii) if $\widehat{L} \in [L^* + \epsilon'/2, L^* + \epsilon']$, then the next sample split is big.

We refer to $|\widehat{R} - R_{\widehat{L}}| \leq 2\delta$ as the *requirement*. Note that it is satisfied for the starting cluster, $(\widehat{L}_s, \widehat{R}_s)$, by assumption, and because $|R_L - R_{L+\delta'}| \leq \delta$ for all L in $[L_s - \delta', L^* + \epsilon']$, by Lemma 3. We will verify that the requirement remains valid until the sample procedure makes a big split. First, suppose $|\widehat{R} - R_{\widehat{L}}| \leq \delta$. By properties (i)-(ii), together with the bound $|R_{\widehat{L}} - R_{\widehat{L}+\delta'}| \leq \delta$, either the next sample split is big or it reduces $(\widehat{L}, \widehat{R})$ to a cluster that still satisfies the requirement. Now consider the case $|\widehat{R} - R_{\widehat{L}}| \in [\delta, 2\delta]$. Consider the last two inequalities in Lemma 3, and, again, derive the sample sample counterparts,

$$\max_{[R-\delta, R]} \widehat{G}_{L,R} \quad > \quad \max_{[L, L+\delta']} \widehat{G}_{L,R} \quad \text{for} \quad L \in [L_s - \delta', L^* + \epsilon'], \quad R \in [R_L + \delta', R_L + 2\delta]$$

$$\max_{[L, L+\delta']} \widehat{G}_{L,R} \quad > \quad \max_{[R-\delta, R]} \widehat{G}_{L,R} \quad \text{for} \quad L \in [L_s - \delta', L^* + \epsilon'], \quad R \in [R_L - 2\delta', R_L - \delta].$$

Taking inequality $|R_{\widehat{L}} - R_{\widehat{L}+\delta'}| \leq \delta$ into account, we deduce that the next sample split, if small, reduces the distance $|\widehat{R} - R_{\widehat{L}}|$. Thus, when the sample procedure is applied to $(\widehat{L}_s, \widehat{R}_s)$, the requirement remains satisfied until the big split, and properties (i)-(iii) remain valid. Consequently, we can use these properties repeatedly to establish that, for all sufficiently large n, the sample procedure starting at $(\widehat{L}_s, \widehat{R}_s)$ makes a number of small splits, followed by a big split, $\widehat{sp} = (\widehat{L}^*, \widehat{s}^*, \widehat{R}^*)$, such that $||\widehat{sp} - sp||_\infty < \epsilon$.

Now consider the case where the population procedure truncates, in a two-sided fashion, the cluster (L_s, R_s) all the way to an empty set. To establish that the sample procedure, applied to $(\widehat{L}_s, \widehat{R}_s)$, makes only small splits, we can apply the same reasoning as in the previous case, but using only the first of the three inequality sets in display (S3) and only the first of the properties (i)-(iii). The rest of the cases can be handled analogously, with only minor modifications to the original argument. Thus, for all sufficiently large n, the sample procedure starting at $(\widehat{L}_s, \widehat{R}_s)$ either makes small splits until the end or makes small splits

until a split $\widehat{sp}$, such that $||\widehat{sp} - sp||_\infty < \epsilon$. Note that the width of the smaller sub-cluster produced by a small split is less than ϵ. Because ϵ can be chosen arbitrarily small, we conclude that $\widehat{sp}$ converges to sp, while the maximum size of the splits prior to $\widehat{sp}$ goes to zero. Note that $size(\widehat{sp})$ converges to $size(sp) = \alpha^*$, by continuity of the function $size$. Then, by the classical Glivenko-Cantelli theorem, $\widehat{size}(\widehat{sp})$ converges to α^* as well.

Because the number of population splits is finite, we can repeat all of the preceding arguments sequentially. For example, to analyze the clustering procedures between the first and, potentially, second population split, we first consider the left sub-cluster and set L_s, R_s, $\widehat{L}_s$ and $\widehat{R}_s$ to L^*, s^*, $\widehat{L}^*$ and $\widehat{s}^*$, respectively. Note that (L_s, R_s) is appropriately bounded, and, for every positive δ and δ', inequalities $|\widehat{R}_s - R_s| \leq \delta$ and $|\widehat{L}_s - L_s| \leq \delta'$ hold for all sufficiently large n. This verifies the assumptions stated in the first paragraph of the proof, which are the only ones needed for the preceding argument. The same is true for the right sub-cluster, on which we set L_s, R_s, $\widehat{L}_s$ and $\widehat{R}_s$ to s^*, R^*, $\widehat{s}^*$ and $\widehat{R}^*$, respectively.

To complete the proof of the theorem, it is only left to show that the maximum size of all the sample splits leading to the very first starting cluster, $(\widehat{L}_s, \widehat{R}_s)$, converges to zero. Recall that the existence of $(\widehat{L}_s, \widehat{R}_s)$, with the required properties, is guaranteed by Lemma 1. The same result also tells us that $P_{-\infty,\widehat{L}_s}$ and $P_{\widehat{R}_s,\infty}$ are bounded above by ϵ. As ϵ can be chosen arbitrarily small, we have established the desired convergence.

Proof of Lemma 1

The proof takes advantage of several applications of Lemma 2 and the strong law of large numbers. We restrict our attention to the set of probability one that remains after casting out the negligible sets associated with the aforementioned convergence results. We conduct a pointwise argument, at a fixed ω, but suppress the dependence on ω to simplify the notation.

When (L_0, R_0), the support of the underlying distribution, is bounded, the conclusion of the lemma holds for $L = L_0$ and $R = R_0$. Consider an unbounded case of $R_0 = \infty$ and $|L_0| < \infty$. Note that the population procedure starts with a right-sided truncation. Let T^* denote the right endpoint of the cluster at which the population procedure either transitions to a two-sided truncation or makes a split. Fix a K_1, such that $P_{K_1,\infty} \leq \epsilon$, with a further requirement that $K_1 > T^*$, if T^* exists. Take K_2 large enough to satisfy $K_2 > \mu_{L_0,\infty} + \mu_{K_1,\infty} - L_0$. By the law of large numbers, $K_2 - \widehat{\mu}_{L_0,\infty} > \widehat{\mu}_{K_1,\infty} - L_0$, for all sufficiently large n. Suppose that n is large enough for the above lower bound to be

satisfied. Then, the following inequalities hold for all $R > K_2$ and $r \in [L_0, K_1]$:

$$\widehat{G}_{L_0,R}(R) = R - \widehat{\mu}_{L_0,R} > K_2 - \widehat{\mu}_{L_0,\infty} > \widehat{\mu}_{K_1,\infty} - L_0 > \widehat{\mu}_{r,R} - \widehat{\mu}_{L_0,r} = \widehat{G}_{L_0,R}(r).$$

This guarantees that for every sample cluster $(L_0, \widehat{R})$, with $\widehat{R} > K_2$, the next sample split point is placed to the right of K_1, which implies that the sample procedure will eventually produce a cluster $(L_0, \widehat{R})$ with $\widehat{R} \in [K_1, K_2]$. Because the population procedure reduces (L_0, ∞) to (L_0, K_1) using the right-sided truncation, interval (L_0, R) is a bounded population cluster for each R in $[K_1, K_2]$. This completes the proof for the case $R_0 = \infty$, $|L_0| < \infty$. The case $L_0 = -\infty$, $R < \infty$ can be handled analogously.

Now consider the case $(L_0, R_0) = (-\infty, \infty)$. Let K_1^- and K_1^+ satisfy $P_{-\infty,K_1^-} \leq \epsilon$ and $P_{K_1^+,\infty} \leq \epsilon$. If the population procedure makes a split, and the first one is applied to the cluster (L^*, R^*), we further require $(K_1^-, K_1^+) \supset [L^*, R^*]$. Additional conditions are placed on K_1^- and K_1^+ below. Take K_2^+ sufficiently large to ensure $K_2^+ > \mu_{K_1^-,\infty} + \mu_{K_1^+,\infty} - \mu_{-\infty,K_1^-}$. By the law of large numbers, $K_2 - \widehat{\mu}_{K_1^-,\infty} > \widehat{\mu}_{K_1^+,\infty} - \widehat{\mu}_{-\infty,K_1^-}$ for all sufficiently large n. Consequently, for all $r \in [K_1^-, K_1^+]$, $R > K_2^+$, $L \leq K_1^-$ we have:

$$\widehat{G}_{L,R}(R) = R - \widehat{\mu}_{L,R} > K_2^+ - \widehat{\mu}_{K_1^-,\infty} > \widehat{\mu}_{K_1^+,\infty} - \widehat{\mu}_{-\infty,K_1^-} \geq \widehat{\mu}_{r,R} - \widehat{\mu}_{L,r} = \widehat{G}_{L,R}(r).$$

This guarantees that for every sample cluster $(\widehat{L}, \widehat{R})$, with $\widehat{L} \leq K_1^-$ and $\widehat{R} > K_2^+$, the next sample split point is placed outside of $[K_1^-, K_1^+]$.

If K_2^- lies below $\mu_{-\infty,K_1^+} + \mu_{-\infty,K_1^-} - \mu_{K_1^+,\infty}$, then, by similar arguments, $\widehat{G}_{L,R}(L) > \widehat{G}_{L,R}(r)$ for all $r \in [K_1^-, K_1^+]$, $L < K_2^-$ and $R \geq K_1^+$. This guarantees that for every sample cluster $(\widehat{L}, \widehat{R})$, with $\widehat{L} < K_2^-$ and $\widehat{R} \geq K_1^+$, the next sample split point is placed outside of $[K_1^-, K_1^+]$.

Consequently, the sample procedure produces a cluster $(\widehat{L}, \widehat{R})$ with $\widehat{L} \in [K_2^-, K_1^-]$ and $\widehat{R} \in [K_1^+, K_2^+]$. Because $K_1^- < L^*$, interval $(\widehat{L}, R_{\widehat{L}})$ is a population cluster. It is only left to show that we can find an $\widehat{R}$ that is located within δ of $R_{\widehat{L}}$. We will use the following result, which is proved further below.

Lemma 4. *Let (L, R_L) be a population cluster, achieved via a two-sided truncation of the real line. If K_1^- is sufficiently small, then for each positive δ we have*

$$\max_{[R_L+\delta,R]} G_{L,R} > \max_{[L,R_L]} G_{L,R} \qquad \textit{for all } L \in [K_2^-, K_1^-] \textit{ and } R \in [R_L + \delta, K_2^+].$$

Applying the argument in the paragraph below the statement of Lemma 3, and replacing L with $\widehat{L}$, we deduce that $\max_{[R_{\widehat{L}}+\delta,R]} \widehat{G}_{\widehat{L},R} > \max_{[\widehat{L},R_{\widehat{L}}]} \widehat{G}_{\widehat{L},R}$ for all $R \in [R_{\widehat{L}}+\delta, K_2^+]$, as long as n is sufficiently large. Hence, if $\widehat{R} > R_{\widehat{L}} + \delta$, the sample procedure sequentially moves the right endpoint of cluster $(\widehat{L}, \widehat{R})$ further left, until it falls in $[R_{\widehat{L}}, R_{\widehat{L}} + \delta]$. The case $\widehat{R} < R_{\widehat{L}} - \delta$ can be handled using analogous arguments.

Proof of Lemma 2

Let P_n denote the empirical measure associated with observations $x_1, ..., x_n$, and let P be the corresponding population distribution. To simplify the presentation, we replace expressions $\int h(x)dP_n(x)$ and $\int h(x)dP(x)$ by P_nh and Ph, respectively.

The classical Glivenko-Cantelli theorem gives the following uniform convergence,

$$\sup_{l<r} |P_n(l,r) - P(l,r)| \to 0, \tag{S4}$$

almost surely, as n goes to infinity. Note that the collection of functions $h_{l,r}(x) = x1_{\{l<x<r\}}$, defined for all real l and r, forms a VC class with an integrable envelope, $H(x) = |x|$. Consequently, by a functional generalization of the Glivenko-Cantelli theorem (e.g. Theorem 19.13 in van der Vaart 1998), we have

$$\sup_{l,r} |P_nh_{l,r} - Ph_{l,r}| \to 0, \tag{S5}$$

almost surely, as n goes to infinity. For the rest of the proof we cast out the negligible sets on which (S4) and (S5) break down, and conduct a pointwise argument on the remaining set of probability one. We suppress the dependence on ω to simplify the notation.

Fix an arbitrarily small positive ϵ. Define $A_M(\epsilon) = A_M \cap \{(l,r) : \ r - l \geq \epsilon\}$. Let $c_\epsilon = \inf_{A_M(\epsilon)} P(l,r)$, and note that c_ϵ is positive. Taking advantage of the convergence in (S4) and (S5), we can bound $\sup_{A_M(\epsilon)} |\widehat{\mu}_{l,r} - \mu_{l,r}|$ above by

$$\begin{aligned}
& \sup_{A_M(\epsilon)} \left| \frac{P_nh_{l,r} - Ph_{l,r}}{P(l,r)} \right| + \sup_{A_M(\epsilon)} |P_nh_{l,r}| \left| \frac{1}{P_n(l,r)} - \frac{1}{P(l,r)} \right| \\
\leq \ & c_\epsilon^{-1} \sup_{l,r} |P_nh_{l,r} - Ph_{l,r}| + c_\epsilon^{-1} \frac{\sup_{l,r} P_n|h_{l,r}|}{\inf_{A_M(\epsilon)} P_n(l,r)} \sup_{l<r} |P_n(l,r) - P(l,r)| \\
= \ & c_\epsilon^{-1} o(1) + c_\epsilon^{-1} \frac{PH + o(1)}{c_\epsilon + o(1)} o(1) = o(1),
\end{aligned}$$

as n tends to infinity. Consequently,

$$\sup_{(L,R)\in A_M} \max_{a\in[L,R]} |\widehat{G}_{L,R}(a) - G_{L,R}(a)| \leq 2\epsilon + 2 \sup_{A_M(\epsilon)} |\widehat{\mu}_{l,r} - \mu_{l,r}| = 2\epsilon + o(1).$$

This completes the proof, because ϵ can be chosen arbitrarily small.

Proof of Lemma 3

For concreteness we assume $L_s > L_0$ throughout the proof. The case $L_s = L_0$ follows with only minor modifications. Suppose that positive ϵ' and δ' are small enough for R_L to be defined on $[L_s - \delta', L^* + \epsilon']$. Then, $\max_{L\in[L_s-\delta',L^*+\epsilon']} |R_L - R_{L+\delta'}| \leq \delta$ for every positive δ, as long as δ' is sufficiently small, by uniform continuity of R_L on a compact set.

Differentiating $G_{L,R}(s)$ with respect to s, we derive the following formulas:

$$G'_{L,R}(L) = \frac{f(L)(\mu_{L,R} - L)}{P_{L,R}} - \frac{1}{2} \quad \text{and} \quad G'_{L,R} = \frac{1}{2} - \frac{f(R)(R - \mu_{L,R})}{P_{L,R}}. \tag{S6}$$

Note that for $L \in [L_s, L^*]$, function $G_{L,R}$ is maximized at the endpoints. Consequently, $G'_{L,R_L}(L) \leq 0$ and $G'_{L,R_L}(R_L) \geq 0$, when $L \in [L_s, L^*]$. For the clarity of the exposition we first focus on the *regular* case, where both inequalities are strict for $L \in [L_s, L^*]$. By continuity of $G'_{L,R}$, there exist positive δ_1 and ϵ_1, such that $G'_{L,R}(L) < 0$ and $G'_{L,R}(R) > 0$ for all $L \in [L_s - \delta_1, L^* + \epsilon_1]$ and $|R - R_L| \leq \delta_1$. Using uniform continuity of $G'_{L,R}$ on compact sets, we conclude that there exists a positive δ_3, such that

$$G_{L,R}(L) \vee G_{L,R}(R) > G_{L,R}(s) \qquad \text{for all} \quad s \in (L, L + \delta_3] \cup [R - \delta_3, R), \tag{S7}$$

when $L \in [L_s - \delta_1, L^* + \epsilon_1]$ and $|R - R_L| \leq \delta_1$. Because $\arg\max G_{L,R_L} = \{L, R_L\}$ for $L \in [L_s, L^*)$, continuity of $G_{L,R}$ implies that, given an arbitrarily small positive ϵ', we can find a positive δ_2, such that

$$G_{L,R}(L) \vee G_{L,R}(R) > \max_{[L+\delta_3, R-\delta_3]} G_{L,R}, \tag{S8}$$

when $L \in [L_s - \delta_2, L^* - \epsilon']$ and $|R - R_L| \leq \delta_2$. Inequalities (S7) and (S8) yield

$$G_{L,R}(L) \vee G_{L,R}(R) > G_{L,R}(s) \qquad \text{for all} \quad s \in (L, R),$$

when $L \in [L_s - \delta_1 \wedge \delta_2, L^* - \epsilon']$ and $|R - R_L| \leq \delta_1 \wedge \delta_2$. This implies the *first set of inequalities* in Lemma 3.

By a similar argument, for every positive ϵ there exist positive ϵ_2 and δ_4, such that

$$G_{L,R}(L) \vee G_{L,R}(R) > \max_{[L+\delta_3, R-\delta_3] \setminus (s^*-\epsilon, s^*+\epsilon)} G_{L,R}, \tag{S9}$$

when $L \in [L_s - \delta_4, L^* + \epsilon_2]$ and $|R - R_L| \leq \delta_4$. Combining this bound with inequality (S7), we derive the *second set of inequalities* in Lemma 3.

In the proof of Proposition 3 we establish inequalities $\mu_{L^*,s^*} \leq (L^* + s^*)/2$ and $\mu_{s^*,R^*} \geq (s^* + R^*)/2$. In our case they are equalities, because $G_{L^*,R^*}(s^*) = G_{L^*,R^*}(L^*) = G_{L^*,R^*}(R^*)$. Taking this fact into account when solving $G'_{L^*,R^*}(s^*) = 0$ we derive a useful equality, $(R^* - s^*)/P_{s^*,R^*} = (s^* - L^*)/P_{L^*,s^*}$. It follows that

$$f(L^*)(R^* - s^*)/P_{s^*,R^*} < 1 \qquad \text{and} \qquad f(L^*)(s^* - L^*)/P_{L^*,s^*} < 1, \tag{S10}$$

because otherwise $f(L^*)(R^* - L^*)/P_{L^*,R^*} \geq 1$, contradicting the fact that $G'_{L,R_L}(L) < 0$ for $L \in [L_s, L^*]$, which holds in the regular setting that we now focus on.

Consider the function $h(L) = G_{L,R_L}(s^*) - G_{L,R_L}(L)$, and note that $h(L^*) = 0$. Inequalities (S10) guarantee that $h'(L^*) > 0$. Hence, there exists a positive ϵ_3, such that for every $\epsilon' \leq \epsilon_3$,

$$G_{L,R_L}(s^*) > G_{L,R_L}(L) = G_{L,R_L}(R_L),$$

when $L \in [L^* + \epsilon'/2, L^* + \epsilon']$. Continuity of $G_{L,R}(s)$, together with compactness of the intervals involved, implies that for every $\epsilon' \leq \epsilon_3$ there exists a positive δ_5, such that

$$G_{L,R}(s^*) > G_{L,R}(L) \vee G_{L,R}(R), \tag{S11}$$

when $L \in [L^* + \epsilon'/2, L^* + \epsilon']$ and $|R - R_L| \leq \delta_5$. Because $\arg\max G_{L^*,R_{L^*}} = \{L^*, R_{L^*}, s^*\}$, we can also find positive ϵ_4 and δ_6, for which

$$G_{L,R}(s^*) > \max_{[L+\delta_3, R-\delta_3] \setminus (s^*-\epsilon, s^*+\epsilon)} G_{L,R}, \tag{S12}$$

when $L \in [L^*, L^* + \epsilon_4]$ and $|R - R_L| \leq \delta_6$. Combining bounds (S7), (S11) and (S12), we establish the *third set of inequalities* in Lemma 3.

Consider function $h_L(R) = G_{L,R}(R) - G_{L,R}(L)$. Differentiation gives $h'_L(R_L) = 1 - 2f(R_L)(R_L - \mu_{L,R_L})/P_{L,R_L}$, which is strictly positive for $L \in [L_s, L^*]$ in the regular case that we now focus on. Taking advantage of equality $h_L(R_L) = 0$, continuity of h'_L and compactness of the intervals involved, we deduce that, provided ϵ', δ and δ' are sufficiently small, inequalities $G_{L,R}(R) > \max_{[L,L+\delta']} G_{L,R}$ are satisfied for all $L \in [L_s - \delta', L^* + \epsilon']$ and $R \in [R_L + \delta, R_L + 2\delta]$. This establishes the *fourth set of inequalities* in Lemma 3. The *fifth set* can be derived using analogous arguments.

We now examine the irregular setting, where $G'_{L,R_L}(L) = 0$ or $G'_{L,R_L}(R_L) = 0$ for some L in $[L_s, L^*]$. For concreteness, we examine the case $G'_{L^i,R^i}(R^i) = 0 \neq G'_{L^i,R^i}(L^i)$, where $L^i \in [L_s, L^*]$ and $R^i = R_{L^i}$. Suppose that $R^i < R_s$; the case $R^i = R_0$ can be handled with minor modifications. Differentiation gives $G''_{L^i,R^i}(R^i) = -f'(R^i)/[6f(R^i)]$, which, together with $\arg\max G_{L^i,R^i} \supseteq \{L^i, R^i\}$, implies $f'(R^i) \geq 0$. However, a strict inequality, $f'(R^i) > 0$ implies that $2f(R_L)(R_L - \mu_{L,R_L}) > P_{L,R_L}$ for all L that are sufficiently close to L^i and satisfy $L < L^i$. Because this contradicts inequality $G'_{L,R_L}(R_L) \geq 0$, we conclude that $f'(R^i) = 0$, and thus, $G'_{L^i,R^i}(R^i) = G''_{L^i,R^i}(R^i) = 0$. Further differentiation yields $G'''_{L^i,R^i}(R^i) = -f''(R^i)/[2f(R^i)]$, which implies $f''(R^i) \leq 0$. This stepwise argument can be continued if $f''(R^i) = 0$, however we will focus on the case $f''(R^i) < 0$, for concreteness. Note that R^i is an interior mode of f.

Let $l = L - L^i$, $t = s - R^i$ and $r = R - R^i$. After deriving the third order Taylor series expansion for the function $(L, s, R) \mapsto G_{L,R}(s)$ at (L^i, R^i, R^i), we establish the following approximation, which holds for (l, t, r) near zero:

$$G_{L^i+l,R^i+r}(R^i + t) - G_{L^i+l,R^i+r}(R^i) = \frac{-f''(R^i)}{6f(R^i)}(3t^3 - t^2 r + 3tr^2) + o(|l|^3 + |t|^3 + |r|^3).$$

Analysis of the above expression reveals that, because $f''(R^i) < 0$, the approximating cubic function is increasing in t for every fixed r and l.

We now revisit the first set of inequalities in Lemma 3; the rest of the sets can be established using similar arguments. For the corresponding proof given in the regular setting to still go through, it is sufficient to establish that for every small enough positive δ_1,

$$G_{L,R}(L) \vee G_{L,R}(R) > G_{L,R}(s) \qquad \text{for all} \quad s \in (L, R - \delta_1/4],$$

when $|L - L^i| \leq \delta_1$ and $|R - R^i| \leq \delta_1$.

By the established monotonicity of the cubic approximation, we can find a positive δ_3, such that for each positive $\delta < 4\delta_3$,

$$G_{L,R}(R) > G_{L,R}(s) \qquad \text{for all} \quad s \in [R - \delta_3, R - \delta_1/4], \tag{S13}$$

when $|L - L^i| \leq \delta_1$ and $|R - R^i| \leq \delta_1$. Because $\arg\max G_{L^i,R^i} = \{L^i, R^i\}$, we can also find a positive δ_2, such that

$$G_{L,R}(R) \vee G_{L,R}(L) > \max_{[L+\delta_3, R-\delta_3]} G_{L,R}, \tag{S14}$$

when $|L - L^i| \leq \delta_2$ and $|R - R^i| \leq \delta_2$. Combining inequalities (S13) and (S14) establishes

$$G_{L,R}(R) \vee G_{L,R}(L) > \max_{[L+\delta_3, R-\delta_1/4]} G_{L,R},$$

when $|L - L^i| \leq \delta_1$, $|R - R^i| \leq \delta_1$, and δ_1 is sufficiently small. Using the regular case argument in the beginning of the proof we also establish the left side bound,

$$G_{L,R}(L) > G_{L,R}(s) \qquad \text{for all} \quad s \in (L, L + \delta_3],$$

when $|L - L^i| \leq \delta_2$, $|R - R^i| \leq \delta_2$, and δ_2 is sufficiently small, which completes the proof.

Proof of Lemma 4

Let R^m be the maximum of R^* and the right-most mode of f. Let L^m be the minimum of L^* and the left-most mode of f. As we point out in the proof of Proposition 3, any interior maximizer of $G_{L,R}$ must lie in (L^m, R^m). Thus, it is sufficient to verify

(i) $G_{L,R}(R) > \max_{[L^m,R^m]} G_{L,R}$ and

(ii) $G_{L,R}(R) > G_{L,R}(L)$,

for all $L \in [K_2^-, K_1^-]$ and $R \in [R_L + \delta, K_2^+]$. If we choose K_1^- sufficiently small, so that $K_1^- < L^m$, $R_{K_1^-} > R^m$ and $R_{K_1^-} > \mu_{-\infty,R_m} + \mu_{R^m,\infty} - \mu_{-\infty,L^m}$, then (i) follows from

$$G_{L,R}(R) > R_{K_1^-} - \mu_{-\infty,R_m} > \mu_{R^m,\infty} - \mu_{-\infty,L^m} > \max_{[L^m,R^m]} G_{L,R}.$$

The existence argument in the proof of Proposition 5 implies that that the derivative of the function $h_L(R) = G_{L,R}(R) - G_{L,R}(L)$ is strictly positive on $[R_L, \infty)$, for every $L < L^m$. Statement (ii) then follows from $h_L(R_L) = 0$.

B Proof of Theorem 2

Note that for the purposes of this proof we can replace "sample frequency" with "underlying probability" in the definition of the set $\widehat{\mathcal{S}}(\tau)$. By the classical Glivenko-Cantelli theorem, the corresponding modified statement of the theorem implies the original one. Write $\mathcal{J}(\tau)$ for the collection of all intervals (L, R), such that such that $P(L, R) \geq \tau$, $P(-\infty, L) \geq \tau$ and $P(R, \infty) \geq \tau$. We will only use the last two inequalities to ensure that, with probability tending to one, all (L, R) are contained in a fixed bounded interval, on which f is bounded away from zero.

Main Body of the Proof

The following two lemmas give us appropriate control over the difference between small perturbations in $\widehat{G}_{L,R}$ and $G_{L,R}$.

Lemma 5. *For every positive ϵ and τ, there exists a $O_p(1)$ random sequence M_n, such that*

$$\left|(\widehat{G}_{L,R}(R) - \widehat{G}_{L,R}(a)) - (G_{L,R}(R) - G_{L,R}(a))\right| \;\leq\; \epsilon(R - a) + (\log n/n)M_n$$

$$\left|(\widehat{G}_{L,R}(L) - \widehat{G}_{L,R}(a)) - (G_{L,R}(L) - G_{L,R}(a))\right| \;\leq\; \epsilon(a - L) + (\log n/n)M_n,$$

for all $(L, R) \in \mathcal{J}(\tau)$ and $a \in (L, R)$, such that $\widehat{G}_{L,R}(a)$ is well-defined.

The proof of Lemma 5 is given further below. The next result can be established using a similar type of argument, however, it is a direct corollary of Lemma 4.1 in Kim and Pollard (1990). Note that the $\log n$ factor, which appears in Lemma 5 but not in Lemma 6, is due to the fact that points L and R, near which function $\widehat{G}_{L,R}$ is approximated in Lemma 5, are allowed to vary, while s, the corresponding point in Lemma 6 is fixed.

Lemma 6. *Given positive τ and ϵ, and a point $s \in (L_0, R_0)$, there exists a $O_p(1)$ sequence M_n, such that*

$$\left|(\widehat{G}_{L,R}(s) - \widehat{G}_{L,R}(a)) - (G_{L,R}(s) - G_{L,R}(a))\right| \leq \epsilon|s-a|^2 + n^{-2/3}M_n,$$

for all L,a and R, such that $(L,s) \in \mathcal{J}(\tau)$, $(s,R) \in \mathcal{J}(\tau)$, $a \in (L,R)$ and $\widehat{G}_{L,R}(a)$ is well-defined.

We start by deriving the uniform rate of convergence for the *small* sample splits. As in the proof of consistency, we focus on the case of the population procedure performing a two-sided truncation, followed by a split. The rest of the cases can be handled using analogous arguments. Let (L^*, s^*, R^*) be the first population split. The population procedure truncates (L_0, R_0), along (L, R_L), down to (L^*, R^*). As we showed in the consistency proof, the sample procedure reduces the cluster consisting of all the observations down to, approximately, (L^*, R^*), by repeatedly splitting off small clusters near the boundary. The split point for the big split is placed near s^*. Given a positive δ, we define $\mathcal{A}(\delta) = \{(L,R) : L \leq L^* + \delta, |R - R_L| \leq \delta\}$. In the consistency proof we showed that the following property holds with probability tending to one for the starting cluster, (L,R), of each small sample split: $(L,R) \in \mathcal{A}(\delta)$ and $\arg\max \widehat{G}_{L,R} \subset [L, L+\delta] \cup [R-\delta, R]$. Let $\widehat{a}_{L,R} = \arg\max_{[L,L+\delta]\cup[R-\delta,R]} \widehat{G}_{L,R}$. Then, to establish bound (10) in the statement of Theorem 2 it is sufficient to prove that for every positive τ there exists a positive δ and a $O_p(\log n/n)$ sequence B_n, such that inequality $(R - \widehat{a}_{L,R}) \wedge (\widehat{a}_{L,R} - L) \leq B_n$ holds for all $(L,R) \in \mathcal{A}(\delta) \cap \mathcal{J}(\tau)$.

To characterize small perturbations in the population criterion function, $G_{L,R}$, we will use the derivations in the proof of Lemma 3 (we focus only on the *regular* part of the proof, because of the new regularity condition, C3). Those derivations imply that there exist positive c_0 and δ, such that, for all $(L,R) \in \mathcal{A}(\delta) \cap \mathcal{J}(\tau)$,

$$G_{L,R}(R) - G_{L,R}(a) \;\geq\; c_0(R-a), \qquad \text{for } a \in [R-\delta, R], \text{ and} \tag{S15}$$

$$G_{L,R}(L) - G_{L,R}(a) \;\geq\; c_0(a-L), \qquad \text{for } a \in [L, L+\delta]. \tag{S16}$$

Define

$$\Delta_n(L,a,R) = \left|(\widehat{G}_{L,R}(R) - \widehat{G}_{L,R}(a)) - (G_{L,R}(R) - G_{L,R}(a))\right| + \left|(\widehat{G}_{L,R}(L) - \widehat{G}_{L,R}(a)) - (G_{L,R}(L) - G_{L,R}(a))\right|.$$

Combining Lemma 5, applied for $\epsilon = c_0/2$, with inequalities (S15) and (S16), we deduce

$$\begin{aligned} 0 &\geq \widehat{G}_{L,R}(R) - \widehat{G}_{L,R}(\widehat{a}_{L,R}) \\ &\geq c_0(R - \widehat{a}_{L,R}) \wedge (\widehat{a}_{L,R} - L) - \Delta_n(L, \widehat{a}_{L,R}, R) \\ &\geq (c_0/2)(R - \widehat{a}_{L,R}) \wedge (\widehat{a}_{L,R} - L) - (\log n/n)M_n. \end{aligned}$$

Hence, $(R - \widehat{a}_{L,R}) \wedge (\widehat{a}_{L,R} - L) \leq (c_0/2)^{-1}(\log n/n)M_n$ for all $(L,R) \in \mathcal{A}(\delta) \cap \mathcal{J}(\tau)$, which is what we needed to prove. Thus, up to the first big split, the sizes of all the small sample splits are uniformly $O_p(\log n/n)$. In the consistency theorem we showed that, with probability tending to one, the number of the big sample splits equals the (finite) number of the population splits. Consequently, the behavior of the sample procedure after the first big split can be handled by repeating the argument given above. This establishes bound (10) in the statement of Theorem 2.

We now focus on deriving the rate of convergence in (9). Recall that (L^*, s^*, R^*) and $(\widehat{L}^*, \widehat{s}^*, \widehat{R}^*)$ denote the first population split and the first big sample split, respectively. Given a positive δ, define $\mathcal{B}(\delta) = \{(L,s,R) : |L - L^*| \leq \delta, |s - s^*| \leq \delta, |R - R^*| \leq \delta\}$. Note that $(\widehat{L}^*, \widehat{s}^*, \widehat{R}^*) \in \mathcal{B}(\delta)$ with probability tending to one. Derivations in the proof of Lemma 3, together with the assumption $G''_{L^*,R^*}(s^*) \neq 0$, imply that there exist positive constants δ, c_1 and c_2, such that

$$G_{L,R}(s^*) - G_{L,R}(a) \geq c_1|a - s^*|^2 - c_2|a - s^*|\left(|L - L^*| + |R - R^*|\right),$$

for all $(L,s,R) \in \mathcal{B}(\delta)$. Define $\widehat{s} = \arg\max_{[s^*-\delta, s^*+\delta]} \widehat{G}_{\widehat{L}^*,\widehat{R}^*}$ and let

$$D_n = \left|\widehat{G}_{\widehat{L}^*,\widehat{R}^*}(s^*) - \widehat{G}_{\widehat{L}^*,\widehat{R}^*}(\widehat{s}) - G_{\widehat{L}^*,\widehat{R}^*}(s^*) - G_{\widehat{L}^*,\widehat{R}^*}(\widehat{s})\right|.$$

We can handle D_n by applying Lemma 6, with sufficiently small ϵ and τ. It follows that, with probability tending to one,

$$\begin{aligned} 0 &\geq \widehat{G}_{\widehat{L}^*,\widehat{R}^*}(s^*) - \widehat{G}_{\widehat{L}^*,\widehat{R}^*}(\widehat{s}) \\ &\geq G_{\widehat{L}^*,\widehat{R}^*}(s^*) - G_{\widehat{L}^*,\widehat{R}^*}(\widehat{s}) - D_n \\ &\geq c_1|\widehat{s} - s^*|^2 - c_2|\widehat{s} - s^*| \left(|\widehat{L}^* - L^*| + |\widehat{R}^* - R^*| \right) - n^{-2/3} M_n. \end{aligned}$$

Consequently,

$$|\widehat{s} - s^*| = O_p\left(n^{-1/3} + |\widehat{L}^* - L^*| + |\widehat{R}^* - R^*| \right).$$

Thus, to establish the rate in (9) for $(\widehat{L}^*, \widehat{s}^*, \widehat{R}^*)$ it is only left to show $|\widehat{L}^* - L^*| = O_p(n^{-1/3})$ and $|\widehat{R}^* - R^*| = O_p(n^{-1/3})$. In the remainder of the proof we establish the last two stochastic bounds, and then extend the rate of convergence derived for $(\widehat{L}^*, \widehat{s}^*, \widehat{R}^*)$ to the subsequent big sample splits.

We again use the derivations in the proof of Lemma 3, from which it follows that, given a positive τ, there exist positive constants δ, c_1 and c_2, such that

$$G_{L,R}(R) - G_{L,R}(a) \geq c_1(R - R_L) - c_2(a - L),$$

for all $(L, R) \in \mathcal{A}(\delta) \cap \mathcal{J}(\tau)$ and $a \in [L, L + \delta]$. Taking advantage of a maximal inequality for the empirical process indexed by a VC class of functions with a square integrable envelope (e.g. Lemma 19.38 in van der Vaart 1998), Lemma 5, and the already established rate of convergence for the small sample splits, we deduce that there exists a $O_p(n^{-1/2})$ sequence C_n, such that

$$\widehat{G}_{L,R}(R) - \max_{a \in [L, L+\delta]} \widehat{G}_{L,R}(a) \geq c_1(R - R_L) - C_n,$$

for all $(L, R) \in \mathcal{A}(\delta) \cap \mathcal{J}(\tau)$. Consequently, when $R - R_L > c_1^{-1} C_n$, the maximizer of $\widehat{G}_{L,R}$ is at least a δ away from L. Because the split points of the small sample splits are uniformly within $O_p(\log n/n)$ of the cluster boundary, it follows that $R - R_L$ is bounded above by a $2c_1^{-1}C_n$ for all the sample clusters, (L, R), produced up to the first big sample

split. A similar argument that focuses on L instead of R derives a lower bound for $R - R_L$. Thus, $|R - R_L| = O_p(n^{-1/2})$, and it is only left to establish $|\widehat{L}^* - L^*| = O_p(n^{-1/3})$.

It follows from the derivations in the proof of Lemma 3 that there exist positive constants δ, c_1, c_2 and c_3, such that

$$G_{L,R}(s^*) - G_{L,R}(a) \geq c_1(L - L^*) - c_2(R - R_L) - c_3[(a - L) \wedge (R - a)], \qquad \text{(S17)}$$

provided $L \in [L^*, L^* + \delta]$, $|R - R^*| \leq \delta$ and $|a - L| \wedge |a - R| \leq \delta$. Handling the stochastic term the same way we did in the paragraph above, we conclude that there exists a $O_p(n^{-1/2})$ sequence T_n, such that inequality

$$\widehat{G}_{L,R}(s^*) \leq \max_{a \in [L, L+\delta] \cup [R-\delta, R]} \widehat{G}_{L,R}(a)$$

implies $L - L^* \leq T_n$. Note that the same T_n is chosen for all intervals (L, R) satisfying $L \in [L^*, L^* + \delta]$ and $|R - R^*| \leq \delta$. The companion lower bound on $L - L^*$ follows from an analogous argument, which replaces the lower bound in display (S17) with an upper bound. The relationship $|\widehat{L}^* - L^*| = O_p(n^{-1/2})$ then follows directly after taking into consideration the rate of convergence already established for the small sample splits.

This completes the derivation of the rate of convergence for the first big sample split. For the subsequent big splits the same argument can be repeated with the following negligible modification. Due to the fact that endpoints of the clusters resulting from the first big sample split are within $O_p(n^{-1/3})$ of their population counterparts, all the $O_p(n^{-1/2})$ expressions should be replaced with $O_p(n^{-1/3})$.

Proof of Lemma 5

We will only establish the first inequality. The second can be proved by using the same argument with the appropriate adjustment of the notation. For points a that are bounded away from R, the first inequality holds by the functional generalization of the Glivenko-Cantelli theorem (e.g. Theorem 19.13 in van der Vaart 1998). Hence, from here on we only focus on $a \in [(L + R)/2, R]$. We will use "$\lesssim$" to mean that inequality "$\leq$" holds when the right hand side is multiplied by a positive constant, which is chosen independently from the involved parameters, such as n, L and R. Note that the statement of the lemma restricts L, a and R to a bounded interval, on which the infumum of f is positive, while the

supremum is finite. In particular, we have $R - a \lesssim P(a, R)$ and $P(a, R) \lesssim R - a$. Define $d_{l,r}(x) = (r - x)(r - l)^{-1} 1_{\{l<x<r\}}$ and note that

$$\begin{aligned}
|\widehat{\mu}_{a,R} - \mu_{a,R}| &= \left| \frac{P_n d_{a,R}(R-a)}{P_n(a,R)} - \frac{P d_{a,R}(R-a)}{P(a,R)} \right| \\
&\leq |P_n d_{a,R} - P d_{a,R}| \frac{(R-a)}{P(a,R)} + |P_n(a,R) - P(a,R)| \frac{(R - \widehat{\mu}_{a,R})}{P(a,R)} \\
&\lesssim |P_n d_{a,R} - P d_{a,R}| + |P_n(a,R) - P(a,R)| =: E_1 + E_2.
\end{aligned}$$

As before, we let $h_{l,r}(x) = x 1_{l<x<r}$. Observe that $|P_n h_{a,R} - P h_{a,R}| \lesssim E_2$, and stochastic bound $\sup_{l<r}(|P_n h_{l,r} - P h_{l,r}|) + (|P_n(l,r) - P(l,r)|) = o_p(1)$ holds by the functional generalization Glivenko-Cantelli theorem. It follows that

$$\begin{aligned}
&|(\mu_{L,R} - \mu_{L,a}) - (\widehat{\mu}_{L,R} - \widehat{\mu}_{L,a})| \\
&= \left| \frac{P h_{L,a} P(a,R) + P h_{a,R} P(L,a)}{P(L,R)P(L,a)} - \frac{P_n h_{L,a} P_n(a,R) + P_n h_{a,R} P_n(L,a)}{P_n(L,R) P_n(L,a)} \right| \\
&\lesssim E_2 + [(R-a) + E_2]\, o_p(1), \qquad \text{(S18)}
\end{aligned}$$

where the $o_p(1)$ term comes from the functional generalization of the Glivenko-Cantelli theorem and is chosen uniformly for all L, R and a in consideration.

It is only left to show that the bound in the statement of the lemma holds for E_1 and E_2. We start with E_2. We need to show that there exists a $O_p(1)$ sequence of random variables M_n, such that

$$|P_n(l,r) - P(l,r)| \leq \epsilon(r-l) + (\log n/n) M_n, \qquad \text{(S19)}$$

for all (l, r) contained within a given bounded set. Define M_n as the infimum of those values for which the above inequality holds. Write $A_{j,n}$ for the set of intervals (l, r) that satisfy $[2^{j-1} - 1] \log n/n < P(l,r) \leq 2^j \log n/n$. Recall that we have restricted our attention to a bounded interval, on which $\|f\|_\infty$ is finite. In the argument that follows the bounded interval is not explicitly present, however we do use the fact that $c_1 := \|f\|_\infty^{-1}$ is

positive. Observe that

$$\begin{aligned} P(M_n > t) &\leq \sum_{j=1}^{\infty} P(\{\exists (l,r) \in A_{j,n} \text{ s.t. } |P_n(l,r) - P(l,r)| > \epsilon(r-l) + t \log n/n\} \\ &\leq \sum_{j=1}^{\infty} P(\{\sup_{A_{j,n}} |P_n(l,r) - P(l,r)| > (\log n/n)[c_1\epsilon(2^{j-1}-1) + t]\} \\ &\leq \sum_{j=1}^{\infty} \frac{E \sup_{A_{j,n}} |P_n(l,r) - P(l,r)|}{(\log n/n)[c_1\epsilon(2^{j-1}-1) + t]} := \sum_{j=1}^{\infty} S_j. \end{aligned}$$

We will bound the summands in the above expression using a maximal inequality from the empirical process theory. Note that the class of indicator functions for all intervals has polynomial bracketing numbers with respect to $L_2(P)$, and its envelope function is identically equal to one. Consequently, an application of Lemma 19.36 in van der Vaart (1998) gives us the following bound:

$$E \sup_{P(l,r)<\theta^2} |P_n(l,r) - P(l,r)| \lesssim n^{-1/2}\theta[1 + \sqrt{-\log(\theta \wedge 1)}] - \log(\theta \wedge 1)/n. \qquad \text{(S20)}$$

We can bound each S_j by applying inequality (S20) with $\theta^2 = 2^j \log n/n$:

$$S_j \leq \frac{\log n/n}{(\log n/n)[\epsilon(2^{j-1}-1) + t]} = \frac{2^{j/2}}{(2^{j-1}-1)\epsilon + t}.$$

It follows that the sum bounding $P(M_n > t)$ can be made arbitrarily small, uniformly over all n, by increasing t. Consequently, $M_n = O_p(1)$.

The proof of the bound for E_1 is essentially identical to the one above. The class of functions $d_{l,r}$, for l and r in a fixed bounded interval, also has polynomial bracketing numbers with respect to $L_2(P)$, and its envelope function is also identically equal to one. Hence, the argument above works for E_1 after each (l,r) is replaced with $d_{l,r}$.

C Proof of Propositions 1-5

Proposition 1

Let $\lambda_0 = 0$. Suppose that the first merge happens at $\lambda = \lambda_1$, the second at $\lambda = \lambda_2$, and so on. We will first show that, with probability one, values λ_k form an increasing sequence. Consider two merges, $C_3 = C_1 \cup C_2$ and $C = C_3 \cup C_4$. For concreteness, we will focus on the case where cluster C_4 exists at the time of the first merge, and establish

$$\frac{\overline{X}_{C_2} - \overline{X}_{C_1}}{|C_3|} \leq \frac{\overline{X}_{C_4} - \overline{X}_{C_3}}{|C|}. \tag{S21}$$

This will complete the proof because of the continuity of the underlying distribution. The complementary case, where C_4 is formed after the first merge can be analyzed analogously. Suppose that the above inequality does not hold. Then, taking into account representation $|C_3|\overline{X}_{C_3} = |C_1|\overline{X}_{C_1} + |C_2|\overline{X}_{C_2}$, we can derive

$$\begin{aligned}
\overline{X}_{C_4} - \overline{X}_{C_2} &= \overline{X}_{C_4} - \overline{X}_{C_3} - |C_1| \cdot |C_2|^{-1} \cdot (\overline{X}_{C_3} - \overline{X}_{C_1}) \\
&= \overline{X}_{C_4} - \overline{X}_{C_3} - |C_1| \cdot |C_3|^{-1} \cdot (\overline{X}_{C_2} - \overline{X}_{C_1}) \\
&< (\overline{X}_{C_2} - \overline{X}_{C_1}) \left(\frac{|C|}{|C_3|} - \frac{|C_1|}{|C_3|} \right) = \frac{|C_2| + |C_4|}{|C_3|} (\overline{X}_{C_2} - \overline{X}_{C_1}).
\end{aligned}$$

The resulting inequality contradicts the merge $C_3 = C_1 \cup C_2$.

We will now verify that the KKT conditions for optimization problem (2) hold for the solutions produced by Algorithm 1. The KKT conditions are satisfied if there exist β_{ij} with $|\beta_{ij}| \leq 1$ and $\beta_{ij} = -\beta_{ji}$, such that for every i:

$$\alpha_i - x_i + \lambda \sum_{j \neq i, \alpha_j \neq \alpha_i} \text{sign}(\alpha_i - \alpha_j) + \lambda \sum_{j \neq i, \alpha_j = \alpha_i} \beta_{ij} = 0. \tag{S22}$$

Write $C(i)$ for the current cluster containing x_i. Taking into account equations (3), the KKT conditions can be rewritten as follows,

$$\overline{X}_{C(i)} - x_i + \lambda \sum_{j \neq i, j \in C(i)} \beta_{ij} = 0. \tag{S23}$$

We will argue by induction over the number of merges and establish the KKT conditions,

(S23), for all i and for all $\lambda = \lambda_k$. If (S23) holds for a particular tuning parameter value λ_k, it also holds for $\lambda > \lambda_k$, provided the cluster $C(i)$ is not modified, if we shrink all the corresponding β_{ij} by a factor of λ_k/λ.

For $\lambda = \lambda_0$, conditions (S23) hold trivially, with each observation forming its own cluster. Suppose we are able to verify the KKT conditions up to the merge $k-1$. Suppose that the k-th merge, at $\lambda = \lambda_k$, is $C = C_1 \cup C_2$. By the discussion above, conditions (S23) hold at $\lambda = \lambda_k$ for all $i \notin C$, and there exist β_{ij} with $|\beta_{ij}| \leq 1$ and $\beta_{ij} = -\beta_{ji}$, such that

$$\overline{X}_{C_1} - x_i + \lambda_k \sum_{j \neq i, j \in C_1} \beta_{ij} = 0 \quad \text{for all } i \in C_1 \tag{S24}$$

$$\overline{X}_{C_2} - x_i + \lambda_k \sum_{j \neq i, j \in C_2} \beta_{ij} = 0 \quad \text{for all } i \in C_2. \tag{S25}$$

We will set $\beta_{ij} = -1$ for each $i \in C_1$ and $j \in C_2$, and we will keep the remaining β_{ij} values intact. We need to show that

$$\overline{X}_C - x_i + \lambda_k \sum_{j \neq i, j \in C} \beta_{ij} = 0 \quad \text{for all } i \in C. \tag{S26}$$

Consider an $i \in C_1$. Equations (S24) imply $\lambda_k \sum_{j \neq i, j \in C_1} \beta_{ij} = x_i - \overline{X}_{C_1}$. Recall that $\lambda_k = (\overline{X}_{C_2} - \overline{X}_{C_1})/|C|$. It follows that

$$\overline{X}_C - x_i + \lambda_k \sum_{j \neq i, j \in C} \beta_{ij} = \overline{X}_C - \overline{X}_{C_1} + (\overline{X}_{C_1} - \overline{X}_{C_2})|C_2|/|C| = 0, \tag{S27}$$

as required. The argument for $i \in C_2$ is analogous, but uses equations (S25) instead of (S24).

Proposition 2

We will prove the result by induction over the number of merges. For each merge we will establish the following claim: the splitting procedure applied to the last formed cluster matches the merging procedure that formed that cluster. It follows that the sequence of clusters formed by the splitting procedure matches the sequence of clusters formed by the merging algorithm.

The claim for merge number one is trivial: the only possible split exactly matches the

first merge. Suppose that the claim has been established for the first k merges. Let C be the cluster formed by the merge $k+1$, which combines clusters C_1 and C_2, with C_1 being the left one. It is only left to show that the first splitting procedure applied to C produces clusters C_1 and C_2. Consider a possible alternative split: $C = C_3 \cup C_4$, with C_3 being the left cluster that is different from C_1. To verify the claim, we need to establish $\overline{X}_{C_2} - \overline{X}_{C_1} > \overline{X}_{C_4} - \overline{X}_{C_3}$. For concreteness, we will focus on the case $C_3 \subset C_1$. The case $C_1 \subset C_3$ can be handled analogously, taking advantage of the inclusion $C_4 \subset C_2$.

Define $C_5 = C_1 \setminus C_3$. Using representations $\overline{X}_{C_1} = \overline{X}_{C_3}|C_3|/|C_1| + \overline{X}_{C_5}|C_5|/|C_1|$ and $\overline{X}_{C_4} = \overline{X}_{C_5}|C_5|/|C_4| + \overline{X}_{C_2}|C_2|/|C_4|$, we can rewrite the desired inequality as

$$\frac{\overline{X}_{C_2} - \overline{X}_{C_5}}{|C_4|} > \frac{\overline{X}_{C_5} - \overline{X}_{C_3}}{|C_1|}. \tag{S28}$$

Suppose that the cluster C_1 was formed by the merge $C_{11} \cup C_{12}$. The induction claim for merges one through k implies that $\overline{X}_{C_{12}} - \overline{X}_{C_{11}}$ maximizes the corresponding difference of the averages over all partitions of C_1. By the monotonicity of the λ values in the merging algorithm, we have $(\overline{X}_{C_2} - \overline{X}_{C_1})/|C| > (\overline{X}_{C_{12}} - \overline{X}_{C_{11}})/|C_1|$, which yields

$$\frac{\overline{X}_{C_2} - \overline{X}_{C_1}}{|C|} > \frac{\overline{X}_{C_5} - \overline{X}_{C_3}}{|C_1|}. \tag{S29}$$

Consequently, if we can establish that

$$\frac{\overline{X}_{C_2} - \overline{X}_{C_5}}{|C_4|} > \frac{\overline{X}_{C_2} - \overline{X}_{C_1}}{|C|}, \tag{S30}$$

then the required inequality (S28) is satisfied. Using representation $\overline{X}_{C_1} = \overline{X}_{C_3}|C_3|/|C_1| + \overline{X}_{C_5}|C_5|/|C_1|$ we can rewrite (S30) as $(\overline{X}_{C_2} - \overline{X}_{C_1})/|C| > (\overline{X}_{C_5} - \overline{X}_{C_3})/|C_1|$. The last inequality is true by (S29), which completes the proof.

Proposition 3

We will first establish that the population procedure produces no splits. More specifically, we will show that for every s that is an extremum of $G_{L,R}(\cdot)$ on the interval (L, R), we have $G_{L,R}(s) < G_{L,R}(L) \vee G_{L,R}(R)$. First, consider the case where s is less than or equal

to the mode of the density, f, restricted to (L, R). Differentiating function $G_{L,R}$ gives

$$G'_{L,R}(s) = \frac{f(s)P_{L,R}}{P_{L,s}P_{s,R}}[\mu_{L,s} + \mu_{s,R} - \mu_{L,R} - s].$$

Thus, if s is an extremum, then $\mu_{s,R} - \mu_{L,R} = s - \mu_{L,s}$ (see Figure S1).

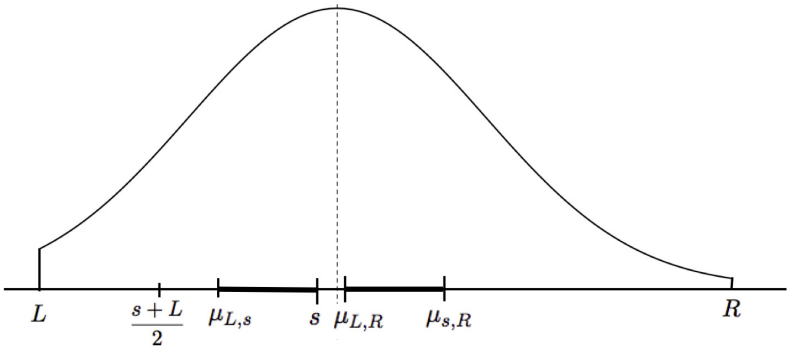

Figure S1: Illustration of the proof of Proposition 3 on a truncated normal density.

Consequently,

$$\begin{aligned} G_{L,R}(s) \;=\; \mu_{s,R} - \mu_{L,s} \;&=\; 2(s - \mu_{L,s}) + \mu_{L,R} - s \\ &<\; s - L + \mu_{L,R} - s = G_{L,R}(L). \end{aligned}$$

The last inequality follows from the fact that f is strictly increasing on (L, s), which implies $s - \mu_{L,s} < (s - L)/2$. Hence, $G_{L,R}(s) < G_{L,R}(L)$. In the case where s is greater than the mode of f on (L, R), an analogous argument establishes $G_{L,R}(s) < G_{L,R}(R)$.

We will now prove that the population clustering procedure is uniquely defined. As we have established in the first part of the proof, the population procedure has the following options: (i) reduce the support to an empty set using a one-sided truncation; (ii) do the same using two-sided truncation; and (iii) start with a one-sided truncation, followed by the two-sided truncation, which reduces the support to an empty set. Because one-sided truncation is always well defined, we can focus on the two-sided truncation exclusively. Consider an interval (L_1, R_1), such that $\arg\max G_{L_1,R_1} = \{L_1, R_1\}$. We need to show that there exists, one and only one, continuous and decreasing function $l \mapsto R_l$, defined on $[L_1, L_2]$, for some $L_2 > L_1$, such that $R_1 = R_{L_1}$ and $L_2 = R_{L_2}$. Recalling the more general existence and uniqueness argument in the proof of Proposition 5, we note that it is

enough to verify that

$$[f(L) \vee f(R)](R - L) < P_{L,R}, \tag{S31}$$

for all $L < R$ with $L \geq L_1$, $R \leq R_1$ and $\arg\max G_{L,R} = \{L, R\}$. Suppose that inequality (S31) is violated. Then, unimodality of f implies $f(L) \neq f(R)$. Another consequence of the violation is that $\mu_{L,R} < (L + R)/2$ if $f(L) > f(R)$ and $\mu_{L,R} > (L + R)/2$, otherwise. This contradicts the equality $\mu_{L,R} = (L + R)/2$, implied by $G_{L,R}(L) = G_{L,R}(R)$. We conclude that $[f(L) \vee f(R)](R - L) < P_{L,R}$. This completes the proof.

Note that in the first part of the proof we establish the following *general property*, which holds for any continuous f: if the population procedure splits (L^*, R^*) at point s^*, then $\mu_{L^*,s^*} \leq (L^* + s^*)/2$ and $\mu_{s^*,R^*} \geq (s^* + R^*)/2$. One consequence of this is that a population split point must lie to the right of left-most mode of f and to the left of the right-most mode.

Proposition 4

Let $\underline{\mathrm{m}}$ denote the interior local minimum of density f. Consider an interval (L, R) that is a subset of the support of f. Suppose there exist points $s_1 < s_2$, such that $\max G_{L,R} = G_{L,R}(s_1) = G_{L,R}(s_2)$. We will establish that for every s in (s_1, s_2) equality $G'_{L,R}(s) = 0$ implies $G''_{L,R}(s) < 0$. Consequently, on the interval (s_1, s_2) function $G_{L,R}$ is not flat and has no local minima, which results in a contradiction.

Define $g_R(s) = f(s)(\mu_{s,R} - s)/P_{s,R}$. Note that

$$G''_{L,R}(s) = f(s)[2g_R(s) - 1]\left(\frac{1}{P_{s,R}} + \frac{1}{P_{L,s}}\right) \quad \text{when} \quad G'_{L,R}(s) = 0. \tag{S32}$$

Because $G'(s_1) = 0$ and $G''_{L,R}(s_1) \leq 0$, we have $2g_R(s_1) - 1 \leq 0$.

In the proof of Proposition 3 we showed that each interior maximizer of $G_{L,R}$ must lie in between the two modes of f. Suppose that s_1 is located to the left of $\underline{\mathrm{m}}$. It follows that f is decreasing on $(s_1, \underline{\mathrm{m}})$. Define $\kappa(s) = P_{s,R}(2g_R(s) - 1)$ and recall that $\kappa(s_1) \leq 0$. Because $\kappa'(s) = 2f'(s)(\mu_{s,R} - s) + f(s)(2g_R(s) - 1)$, we have $\kappa'(s_1) < 0$, and, consequently, $\kappa'(s) < 0$ and $\kappa(s) < 0$ for all s in $(s_1, \underline{\mathrm{m}}]$. Thus, it follows from formula (S32) that, for s in $(s_1, \underline{\mathrm{m}}]$, equality $G'_{L,R}(s) = 0$ implies $G''_{L,R}(s) < 0$.

Define $g_L(s) = f(s)(s - \mu_{L,s})/P_{L,s}$, and note that $g_R(s) = g_L(s)$ when $G'_{L,R}(s) = 0$. An argument analogous to the one in the previous paragraph, but with s_2 and g_L used

instead of s_1 and g_R, establishes that, for s in $(\underline{\mathrm{m}}, s_2)$, equality $G'_{L,R}(s) = 0$ implies $G''_{L,R}(s) < 0$. This completes the proof.

Proposition 5

We will first focus on the case where the two-sided truncation starts at some bounded interval (L_1, R_1). Thus, $R_{L_1} = R_1$. We will show that there exists a point $L^* > L_1$ and a small $\epsilon > 0$, such that the function R_L, with the required properties, is uniquely defined either for $L \in (L_1 - \epsilon, L^*]$, where $L^* = R_{L^*}$, or for $L \in (L_1 - \epsilon, L^* + \epsilon)$, where the population procedure splits cluster (L^*, R_{L^*}). Recall that the required properties of R_L are that it is continuous and decreasing, $\arg\max G_{L,R_L} = \{L, R_L\}$ for $L \in [L_1, L^*)$, $\arg\max G_{L^*,R_{L^*}} \supseteq \{L^*, R_{L^*}\}$, and $G_{L,R_L}(L) = G_{L,R_L}(R_L)$ for $L \in (L_1 - \epsilon, L^* + \epsilon)$.

Write $F(L, R)$ for $G_{L,R}(L)/2 - G_{L,R}(R)/2$, and note that $F(L, R) = \mu_{L,R} - (R+L)/2$. Observe that $\arg\max G_{L,R} = \{L, R\}$ implies $F(L, R) = 0$. To simplify the notation we write $F_1(L, R)$ for the derivative $\partial F(l, r)/\partial l$, evaluated at (L, R), and write $F_2(L, R)$ for the corresponding value of $\partial F(l, r)/\partial r$. Differentiation of F gives us

$$F_1(L, R) = \frac{f(L)(\mu_{L,R} - L)}{P_{L,R}} - \frac{1}{2} \quad \text{and} \quad F_2(L, R) = \frac{f(R)(R - \mu_{L,R})}{P_{L,R}} - \frac{1}{2}. \tag{S33}$$

According to display (S6) in the proof of Lemma 3, we have $F_1(L, R) = G'_{L,R}(L)$ and $F_2(L, R) = -G'_{L,R}(R)$. Because $\arg\max G_{L_1,R_1} = \{L_1, R_1\}$, we then also have inequalities $F_1(L_1, R_1) \leq 0$ and $F_2(L_1, R_1) \leq 0$. Suppose, first, that both of the above inequalities are strict. Then, both $F_1(L_1, R_1)$ and $F_2(L_1, R_1)$ are negative. By the classical implicit function theorem, function $L \mapsto R_L$, for which $F(L, R_L) = 0$, and hence $G_{L,R_L}(L) = G_{L,R_L}(R_L)$, is uniquely defined on $(L_1 - \epsilon, L_1 + \epsilon)$, for some small positive ϵ. By the same theorem R_L is also continuous and decreasing.

As L increases from the value L_1, we define R_L as the largest R below R_1, for which $F(L, R) = 0$. Note that the smallest R_L can be is L. We continuously decrease L until $F_1(L, R_L) = 0$, or $F_2(L, R_L) = 0$, or $L = R_L$, and let L_2 be the value of L corresponding to the stopping point. Note that, by the implicit function theorem, function R_L is continuous and decreasing on $(L_1 - \epsilon, L_2)$, as well as uniquely defined. Because R_L is continuous, decreasing and bounded below by L_2, we can define it at L_2 by continuity, also preserving the monotonicity and the lower bound. We write R_2 for R_{L_2}.

If $L_2 = R_2$ or $L^* < L_2$, then the proof of existence and uniqueness of the two-sided truncation is complete. Otherwise, we need to handle the case where $F_1(L_2, R_2) = 0$ or $F_1(L_2, R_2) = 0$. Note that $G_{L_2,R_2}(L_2) = G_{L_2,R_2}(R_2) = \max G_{L_2,R_2}$. We will focus on the case $F_1(L_2, R_2) = F_2(L_2, R_2) = 0$. The remaining two cases follow from analogous arguments.

The second order Taylor series approximation to $F(L, R)$ at (L_2, R_2) is

$$F(L,R)-F(L_2,R_2) = \frac{(R_2-L_2)}{2P_{L_2,R_2}}\left[f'(L_2)(\Delta L)^2 + f'(R_2)(\Delta R)^2\right] + o\left(|\Delta L|^2 + |\Delta R|^2\right),$$

where $\Delta R = R - R_2$ and $\Delta L = L - L_2$. As we point out in the proof of Lemma 3, equality $G'_{L_2,R_2}(R_2) = F_2(L_2, R_2) = 0$ implies $f'(R_2) \geq 0$. Similarly, we have $f'(L_2) \leq 0$. Suppose, first, that both inequalities are strict; the remaining case can be handled using analogous arguments that use higher order Taylor expansions. Then, given an arbitrarily small neighborhood $\mathcal{N}$ of (L_2, R_2) we can find $(L, R) \in \mathcal{N}$, such that $L > L_2$ and $R > R_2$ and $F(L, R) = 0$. Also note that $F_1(L, R)$ and $F_2(L, R)$ are both negative. Consequently, by the implicit function theorem, the map $l \mapsto R_l$ is uniquely defined in a neighborhood of L, such that $R = R_L$, and the function R_l is continuous and decreasing. Thus, the map is uniquely defined on $(L_2, L_2 + \epsilon)$, for some positive ϵ, it is continuous and decreasing, and $R_l \uparrow R_2$ as $l \downarrow L_2$. Thus, the continuity and monotonicity of the map is preserved when we glue together the functions defined on $(L_1, L_2]$ and $(L_2, L_2 + \epsilon)$. The above Taylor approximation implies that when ϵ is sufficiently small, derivatives $F_1(L, R)$ and $F_2(L, R)$ are both negative for $L \in (L_2, L_2 + \epsilon)$. Thus, we are back in the regular setting of the previous paragraph. Applying the above arguments sequentially establishes that the function R_L is uniquely defined, in a continuous and decreasing fashion, either for $L \in (L_1 - \epsilon, L^*]$, where $L^* = R_{L^*}$, or for $L \in (L_1 - \epsilon, L^* + \epsilon)$, where the population procedure splits cluster (L^*, R_{L^*}). The remaining case $(L_1, R_1) = (-\infty, \infty)$ can be handled using an additional, but similar, argument, where L is decreased, rather than increased from the value L_1.

Thus, we have shown that the two-sided truncation of the population procedure is uniquely defined. Note that the one sided truncation exists and is unique directly by definition, while the uniqueness of the split is ensured by the regularity condition $C2$. It is only left to show that the population procedure will have finitely many splits.

In the proof of Proposition 3 we establish the following general property: a population

split point must lie to the right of left-most mode of the truncated version of f, and to the left of the corresponding right-most mode, where the truncation is done to the cluster that is being split. We assume that this is not the case and conduct an argument to reach a contradiction. The assumption implies that there exists an interval (L^i, R^i), such that for every positive ϵ there exists an infinitely large nested sequence of distinct population clusters, of the form (L, R), sandwiched between $(L^i - \epsilon, R^i + \epsilon)$ and $[L^i, R^i]$, such that either the left or right endpoint of each cluster is a population split point for the previous cluster in the sequence. Without loss of generality, we focus on the case where each right endpoint of the considered cluster sequence is a split point for the previous cluster. By continuity of $G_{L,R}(s)$ we must then have $G_{L^i,R^i}(R^i) \geq G_{L^i,R^i}(L^i)$. Below we examine the case where the last inequality is strict; the case of equality can be handled using an analogous argument. It follows that when ϵ is sufficiently small, the population procedure cannot place a split point in $(L, L^i]$ or do a left sided truncation of the cluster (L, R). Consequently, we have $L = L^i$ for each cluster in the considered sequence.

Recall that $G'_{L,R}(R) = 1/2 - f(R)(R - \mu_{L,R})/P_{L,R}$. If $G'(L^i, R^i)(R^i)$ is negative, then, for a sufficiently small ϵ, $G'_{L,R}(R)$ is negative and bounded away from zero. Consequently, when R gets sufficiently close to R^i, the corresponding split point must lie below R^i, which is a contradiction. Similarly, we reach a contradiction when $G'(L^i, R^i)(R^i)$ is positive. For the rest of the proof we focus on the remaining case $G'(L^i, R^i)(R^i) = 0$.

The general property given in the beginning of the paragraph before the previous one implies that $f'(R) > 0$ for $R \in (R^i, R^i + \epsilon)$ and $f'(R^i) \geq 0$. We first consider the case where the last inequality is strict. Let $t = s - R^i$ and $r = R - R^i$. After deriving the second order Taylor series expansion for the function $(s, R) \mapsto G_{L^i,R}(s)$ at (R^i, R^i), we establish approximation

$$G_{L^i,R^i+r}(R^i + t) - G_{L^i,R^i+r}(R^i) = \frac{-f'(R^i)}{6f(R^i)}(t^2 + tr) + o(|t|^2 + |r|^2),$$

which holds for (t, r) near zero. Because we now focus on the case $f'(R^i) > 0$, we must have $G(L^i, R^i + r)(R^i - r/2) > \max_{s \in [R^i, R^i + r]} G(L^i, R^i + r)(s)$, for every sufficiently small positive r. Consequently, when ϵ is sufficiently small, the population procedure applied to an interval (L^i, R) satisfying $R \in (R^i, R^i + \epsilon)$ must place the split point below R^i. This is again a contradiction of our assumption.

Now consider the remaining case $f'(R^i) = 0$. Let $t = s - R^i$ and $r = R - R^i$.

After deriving the third order Taylor series expansion for the function $(s, R) \mapsto G_{L^i,R}(s)$ at (R^i, R^i), we establish approximation

$$G_{L^i,R^i+r}(R^i + t) - G_{L^i,R^i+r}(R^i) = \frac{-f''(R^i)}{6f(R^i)}(3t^3 - t^2r + 3tr^2) + o(|t|^3 + |r|^3),$$

for (t, r) near zero. We will focus on the case $f''(R^i) \neq 0$, as the case $f''(R^i) = 0$ can be handled using analogous arguments that use higher order Taylor approximation. Because R^i is a local minimum of f, we must have $f''(R^i) > 0$. Analysis of the above Taylor expansion reveals that $G(L^i, R^i + r)(R^i - r) > \max_{s\in[R^i,R^i+r]} G(L^i, R^i + r)(s)$, for every sufficiently small positive r. Consequently, when ϵ is sufficiently small, the population procedure applied to an interval (L^i, R) satisfying $R \in (R^i, R^i + \epsilon)$ must place the split point below R^i. This contradiction completes the proof.

D Analysis of the Population Procedure for Gaussian Mixtures

Consider the case where f is a mixture of two Gaussian densities on the real line. Note that regularity condition C1 is satisfied. Propositions 3 and 4 imply that condition C2 is satisfied as well. Thus, the population procedure is well defined, and, according to the results in Section 3, the sample clustering procedure is consistent. In Table 1 we document the behaviour of the population clustering procedure for various mixtures of two Gaussian distributions on the real line. For 7 different levels of separation between the two normal means we consider 9 different mixing proportions, from the symmetric case of $50 : 50$ mixing to the highly non-symmetric $10 : 90$ mixing. The behavior of the population splitting procedure in other cases can be interpolated from the table using continuity arguments. We present, where applicable, the location of the split point, s^* ("NO" denotes that no splits was detected), as well as the endpoints, L^* and R^*, of the corresponding cluster. The local minimum of the density, $\underline{\mathrm{m}}$, and the split point minimizing the expected misclassification error, s_{MC}, are also provided. We also report, under *Excess MCE*, how much the misclassification error of the population clustering procedure exceeds that of the s_{MC} based oracle rule. Even though the population procedure focusses on cluster separation, rather than classification, its *Excess MCE* is well controlled when the number of splits is detected

correctly. In all cases where s^* is not reported, we checked that the two-sided truncation takes the left endpoint of the cluster all the way to $\underline{\mathrm{m}}$, and thus reduces the density to a unimodal one. Proposition 3 then implies that the population procedure does not produce any splits. Similarly, in the cases where s^* is provided, we verified that the bimodal sub-cluster is truncated down to a unimodal one, which means that the procedure does not produce a second split.

D.1 Additional Technical Details

Note that the population procedure for a distribution supported by the real line starts with a two-sided truncation of the support. The truncation proceeds along the interval (L, R_L), as L is increased, and stops when either the support is reduced all the way to an empty set or the maximum of G_{L,R_L} is achieved at an interior point. In the latter case, define L^* as the smallest L for which there exists an s in (L, R_L), such that $\mu_{L,s} = (L + s)/2$ and $\mu_{s,R_L} = (s + R_L)/2$. Write s^* for the corresponding point s, and let R^* stand for R_{L^*}. The population procedure truncates the support to the interval (L^*, R^*), which is then split at s^*. Note that L^* must be smaller than $\underline{\mathrm{m}}$, the local minimum of the density, by Proposition 3.

Because analytical solutions are not available, we find L^*, R^* and s^* numerically. For each L on a dense grid we locate R_L using equation $\mu_{L,R_L} = (L + R_L)/2$. The symmetry of Gaussian distributions implies that in the search for L^* we only need to consider the values of L that satisfy inequality $L \geq 2(\mu_1 - \mu_2) + R_L$. If the above inequality holds, then we focus on s, the local maximum of G_{L,R_L}, and compute $\delta_1 = \mu_{L,s} - (L + s)/2$ and $\delta_2 = \mu_{s,R_L} - (s + R_L)/2$. The proof of Proposition 5 shows that at $L = L^*$ we have $\delta_1 = \delta_2 = 0$, and, as L crosses L^*, inequalities $\delta_1 < 0$ and $\delta_2 > 0$ are satisfied for the first time. Thus, as L is increased, L^* can be taken as the first point where the above inequalities hold. To check for the second split we apply the same approach on the bimodal sub-cluster. In all the reported cases, we verify that the two-sided truncation takes an endpoint of the sub-cluster all the way to $\underline{\mathrm{m}}$, and thus reduces the density to a unimodal one. Proposition 3 then implies that the population procedure does not produce a second split.

We now provide closed form expressions for $\mu_{L,R}$, $G_{L,R}$ and $G'_{L,R}$. Figure S2 displays $G_{L,R}(a)$, f and $G'_{L,R}$ (normalized), corresponding to the $0.35\,N(-4, 1) + 0.65\,N(4, 1)$ distribution, for three values of L: (a) $L < L^*$, (b) $L = L^*$, (c) $L > L^*$. Locations of $\mu_{L,s}, (L + s)/2, \mu_{s,R_L}$ and $(s + R_L)/2$ are also provided.

CASE	p_1	p_2	μ_1	μ_2	$\underline{\mathrm{m}}$	$s^\star$	$L^\star$	$R^\star$	2nd split	s_{MC}	Excess MCE
	0.50	0.50	-4.50	4.50	0.00	0.00	-8.99	8.99	NO	0.00	0.000
	0.45	0.55	-4.50	4.50	-0.02	-0.45	-8.53	9.43	NO	-0.02	0.000
	0.40	0.60	-4.50	4.50	-0.05	-0.90	-8.08	9.88	NO	-0.04	0.000
	0.35	0.65	-4.50	4.50	-0.07	-1.36	-7.62	10.33	NO	-0.07	0.000
$\|\mu_2 - \mu_1\| = 9$	0.30	0.70	-4.50	4.50	-0.10	-1.82	-7.17	10.79	NO	-0.09	0.001
	0.25	0.75	-4.50	4.50	-0.13	-2.31	-6.67	11.24	NO	-0.12	0.004
	0.20	0.80	-4.50	4.50	-0.16	-2.90	-6.09	11.70	NO	-0.15	0.011
	0.15	0.85	-4.50	4.50	-0.20	-3.82	-5.09	12.16	NO	-0.19	0.037
	0.10	0.90	-4.50	4.50	-0.26	NO				-0.24	0.100
	0.50	0.50	-4.00	4.00	0.00	0.00	-7.99	7.99	NO	0.00	0.000
	0.45	0.55	-4.00	4.00	-0.03	-0.40	-7.58	8.38	NO	-0.03	0.000
	0.40	0.60	-4.00	4.00	-0.05	-0.80	-7.17	8.78	NO	-0.05	0.000
	0.35	0.65	-4.00	4.00	-0.08	-1.22	-6.77	9.19	NO	-0.08	0.001
$\|\mu_2 - \mu_1\| = 8$	0.30	0.70	-4.00	4.00	-0.11	-1.64	-6.34	9.59	NO	-0.11	0.003
	0.25	0.75	-4.00	4.00	-0.15	-2.12	-5.86	9.99	NO	-0.14	0.008
	0.20	0.80	-4.00	4.00	-0.18	-2.72	-5.25	10.40	NO	-0.17	0.020
	0.15	0.85	-4.00	4.00	-0.23	NO				-0.22	0.150
	0.10	0.90	-4.00	4.00	-0.29	NO				-0.28	0.100
	0.50	0.50	-3.50	3.50	0.00	0.00	-6.98	6.98	NO	0.00	0.000
	0.45	0.55	-3.50	3.50	-0.03	-0.35	-6.63	7.34	NO	-0.03	0.000
	0.40	0.60	-3.50	3.50	-0.06	-0.71	-6.27	7.69	NO	-0.06	0.001
	0.35	0.65	-3.50	3.50	-0.10	-1.09	-5.90	8.04	NO	-0.09	0.003
$\|\mu_2 - \mu_1\| = 7$	0.30	0.70	-3.50	3.50	-0.13	-1.49	-5.49	8.39	NO	-0.12	0.007
	0.25	0.75	-3.50	3.50	-0.17	-1.97	-5.01	8.75	NO	-0.16	0.016
	0.20	0.80	-3.50	3.50	-0.22	-2.66	-4.32	9.12	NO	-0.20	0.040
	0.15	0.85	-3.50	3.50	-0.27	NO				-0.25	0.150
	0.10	0.90	-3.50	3.50	-0.34	NO				-0.31	0.100
	0.50	0.50	-3.00	3.00	0.00	0.00	-5.99	5.99	NO	0.00	0.000
	0.45	0.55	-3.00	3.00	-0.04	-0.32	-5.66	6.28	NO	-0.03	0.001
	0.40	0.60	-3.00	3.00	-0.08	-0.64	-5.34	6.59	NO	-0.07	0.004
	0.35	0.65	-3.00	3.00	-0.12	-0.99	-4.99	6.89	NO	-0.10	0.008
$\|\mu_2 - \mu_1\| = 6$	0.30	0.70	-3.00	3.00	-0.16	-1.39	-4.59	7.20	NO	-0.14	0.016
	0.25	0.75	-3.00	3.00	-0.21	-1.91	-4.07	7.52	NO	-0.18	0.034
	0.20	0.80	-3.00	3.00	-0.26	NO				-0.23	0.200
	0.15	0.85	-3.00	3.00	-0.33	NO				-0.29	0.150
	0.10	0.90	-3.00	3.00	-0.41	NO				-0.37	0.100
	0.50	0.50	-2.50	2.50	0.00	0.00	-4.97	4.97	NO	0.00	0.000
	0.45	0.55	-2.50	2.50	-0.05	-0.30	-4.68	5.23	NO	-0.04	0.005
	0.40	0.60	-2.50	2.50	-0.10	-0.61	-4.37	5.49	NO	-0.08	0.011
	0.35	0.65	-2.50	2.50	-0.15	-0.96	-4.01	5.75	NO	-0.12	0.021
$\|\mu_2 - \mu_1\| = 5$	0.30	0.70	-2.50	2.50	-0.20	-1.41	-3.56	6.02	NO	-0.17	0.041
	0.25	0.75	-2.50	2.50	-0.26	NO				-0.22	0.250
	0.20	0.80	-2.50	2.50	-0.33	NO				-0.28	0.200
	0.15	0.85	-2.50	2.50	-0.41	NO				-0.35	0.149
	0.10	0.90	-2.50	2.50	-0.53	NO				-0.44	0.100
	0.50	0.50	-2.00	2.00	0.00	0.00	-3.89	3.89	NO	0.00	0.000
	0.45	0.55	-2.00	2.00	-0.07	-0.32	-3.62	4.15	NO	-0.05	0.015
	0.40	0.60	-2.00	2.00	-0.14	-0.67	-3.28	4.38	NO	-0.10	0.034
	0.35	0.65	-2.00	2.00	-0.21	-1.12	-2.85	4.62	NO	-0.15	0.065
$\|\mu_2 - \mu_1\| = 4$	0.30	0.70	-2.00	2.00	-0.28	NO				-0.21	0.298
	0.25	0.75	-2.00	2.00	-0.37	NO				-0.28	0.248
	0.20	0.80	-2.00	2.00	-0.47	NO				-0.35	0.198
	0.15	0.85	-2.00	2.00	-0.58	NO				-0.43	0.148
	0.10	0.90	-2.00	2.00	-0.74	NO				-0.55	0.097
	0.50	0.50	-1.50	1.50	0.00	0.00	-2.68	2.68	NO	0.00	0.000
	0.45	0.55	-1.50	1.50	-0.12	-0.50	-2.29	2.97	NO	-0.07	0.057
	0.40	0.60	-1.50	1.50	-0.24	NO				-0.14	0.396
	0.35	0.65	-1.50	1.50	-0.38	NO				-0.21	0.344
$\|\mu_2 - \mu_1\| = 3$	0.30	0.70	-1.50	1.50	-0.53	NO				-0.28	0.293
	0.25	0.75	-1.50	1.50	-0.71	NO				-0.37	0.241
	0.20	0.80	-1.50	1.50	-1.50	NO				-0.46	0.190
	0.15	0.85	-1.50	1.50	-1.50	NO				-0.58	0.139
	0.10	0.90	-1.50	1.50	-1.50	NO				-0.73	0.090

Table 1: Finding the population splits for 2-normal mixtures: $p_1\, N(\mu_1, 1) + p_2\, N(\mu_2, 1)$.

If the population density f is a mixture of k normals: $f(x) := \sum_{i=1}^{k} w_i \phi(x - \mu_i)$ then using the following inequality,

$$\int x\, \phi(x-\mu)\, dx = \mu\, \Phi(x-\mu) - \phi(x-\mu),$$

we get closed from expressions of $G_{L,R}$ which can be subsequently optimized:

$$G_{L,R}(a) = P_{a,R}^{-1} \int_a^R x f(x)\, dx - P_{L,a}^{-1} \int_L^a x f(x),\, dx \quad \text{where}$$

$$\int_a^R x f(x)\, dx = \sum_{i=1}^{k} w_i \mu_i \{\Phi(R-\mu_i) - \Phi(a-\mu_i)\} - f(R) + f(a),$$
$$\int_L^a x f(x)\, dx = \sum_{i=1}^{k} w_i \mu_i \{\Phi(a-\mu_i) - \Phi(L-\mu_i)\} - f(a) + f(L), \text{ and}$$
$$P_{a,R} = \sum_{i=1}^{k} w_i \big\{\Phi(R-\mu_i) - \Phi(a-\mu_i)\big\}, \quad P_{L,a} = \sum_{i=1}^{k} w_i \big\{\Phi(a-\mu_i) - \Phi(L-\mu_i)\big\}.$$

The following alternative expression for $G_{L,R}(a)$ in terms of the conditional means and density $f_{L,R}(x) = P_{L,R}^{-1}\, f(x)\, \mathrm{I}\{L, R\}$, where $\tilde{P}_{a_1,a_2} = P_{L,R}^{-1}\, P_{a_1,a_2}$, is also useful:

$$\begin{aligned} G_{L,R}(a) &= \tilde{P}_{L,a}^{-1} \int_a^R x\, f_{L,R}(x)\, dx - \tilde{P}_{a,R}^{-1} \int_L^a x\, f_{L,R}(x)\, dx \\ &= \{\tilde{P}_{L,a}\, \tilde{P}_{a,R}\}^{-1} \left\{ \tilde{P}_{L,a}\, \mu_{L,R} - \int_\infty^a x\, f_{L,R}(x)\, dx \right\} \\ &= f_{L,R}(a) \{\tilde{P}_{L,a}\, \tilde{P}_{a,R}\}^{-2} \left[\mu_{L,R} \tilde{P}_{a,L}^2 + (1 - 2\, \tilde{P}_{a,L}) \int_{-\infty}^a x\, f_{L,R}(x)\, dx - \tilde{P}_{a,L}\, \tilde{P}_{a,R} \right]. \end{aligned}$$

Differentiating the above with respect to a, we arrive at: $G'_{L,R}(a) = \kappa_{L,R}(a) \times H_{L,R}(a)$, where

$$\begin{aligned} \kappa_{L,R}(a) &= f(a)\, P_{L,a}^{-2}\, P_{a,R}^{-2}\, P_{L,R}^{-1}, \quad \text{and} \\ H_{L,R}(a) &= \mu_{L,R} P_{L,a}^2 + \{P_{a,R} - P_{L,a}\} \int_L^a x f(x)\, dx - a\, P_{a,L}\, P_{a,R}. \end{aligned}$$

Note that $\kappa_{L,R}(a) > 0$ for all $a \in (L, R)$. Hence, to track the the extremas of $G_{L,R}$, it is enough to search for the zeros of $H_{L,R}(a)$. We call $H_{L,R}$ the normalized $G'_{L,R}$. Figure S2

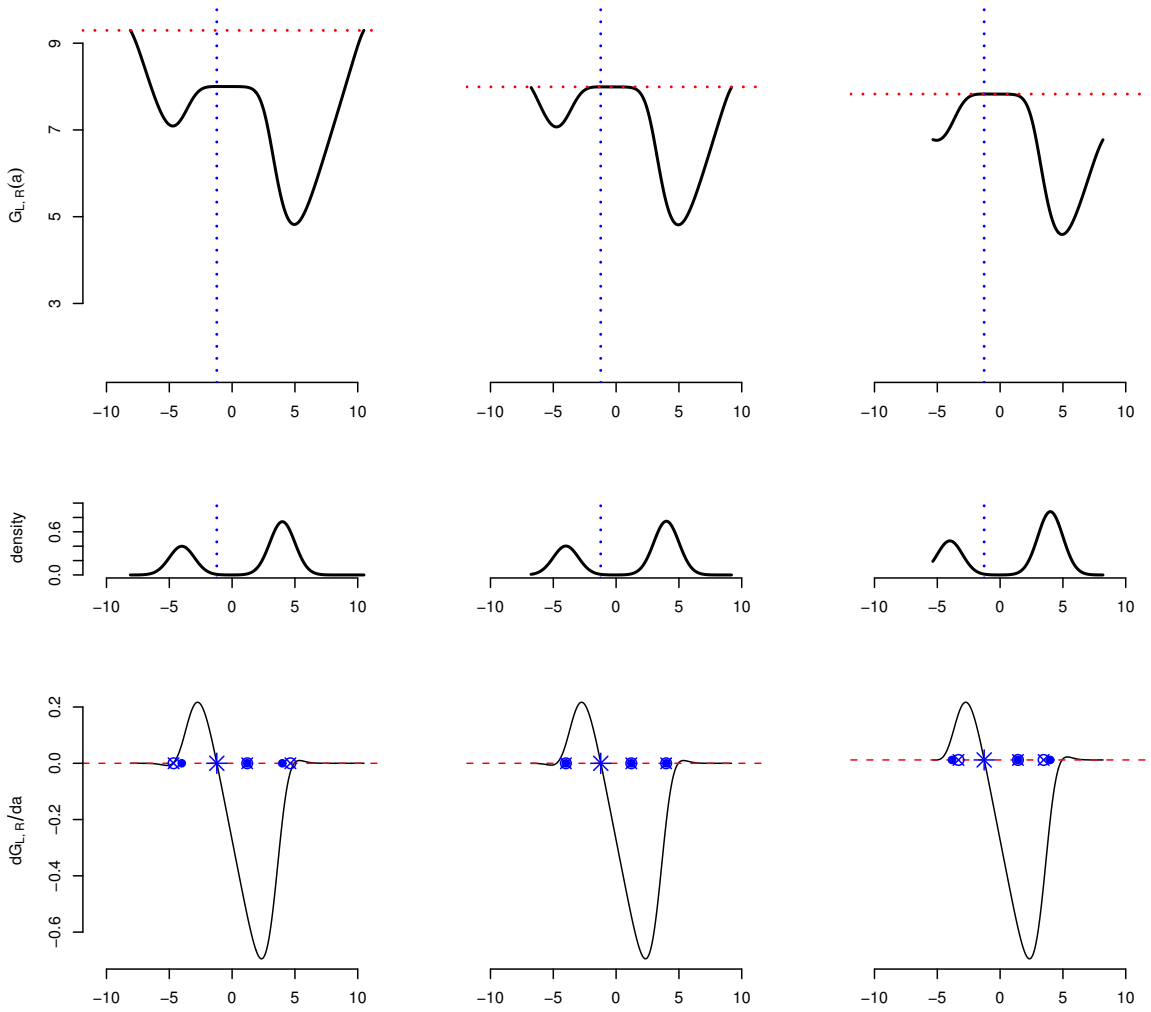

Figure S2: Across rows we have the plots of $G_{L,R}(a)$, the truncated density and the normalized $G'_{L,R}(a)$, as L varies across columns over the 3 cases: (i) $L < L^\star$ (ii) $L = L^\star$ (iii) $L > L^\star$. The population density used here is $0.35\,N(-4,1) + 0.65N\,(4,1)$. The dotted blue line and the stars denote the position of the zero of G'. In the last row, the conditional means, μ_{Ls}, μ_{LR} and μ_{sR}, are denoted by circles and the corresponding mid-points, $(L+s)/2$, $(L+R_L)/2$ and $(s+R_L)/2$, by squares. The signs differences δ_1 and δ_2, defined in Section 3, vary as follows (i) $(\delta_1 > 0, \delta_2 < 0)$ (ii) $(\delta_1 = 0, \delta_2 = 0)$ (iii) $(\delta_1 < 0, \delta_2 > 0)$.

shows the plots of $G_{L,R}(a)$, the truncated density and $H_{L,R}(a)$, when $f = 0.35\,N(-4,1) + 0.65N\,(4,1)$. The plot of $G_{L,R}$ appears flat in the neighborhood of the split. But, the plot of $H_{L,R}$ clearly shows only one zero-crossing and demonstrates uniqueness of the maximum of $G_{L,R}$ in the case of interest.

E Further Details on the Simulation Study & Real Data Analysis

For the numerical experiments in Section 4, we implemented the BMT with threshold α uniformly set at 10%, and with the adjustment of at most 50% truncation for the first split. The details are provided in Algorithm 3.

E.1 Modality Assessment

We compare the performance of the BMT with the following two popular modality assessment procedures:

Silverman Test is based on a kernel density estimate. It uses the idea that if the population density is non-unimodal, a large value of the bandwidth will be required to smooth the data to a unimodal density estimate (Silverman, 1981) . The test uses the minimum bandwidth that produces a unimodal kernel estimator. Large values of the minimum bandwidth based test-statistic provide evidence to support the alternative hypothesis of multi-modality. To conduct the Silverman test, we use the R-package referenced in Vollmer *et al.* (2013). It is based on Gaussian kernels and incorporates Hall and York (2001) adjustment for calculating the p-value.

The Dip Test proposed by Hartigan and Hartigan (1985) is a histogram based method, which does not require estimating the density. The Dip-statistic is the minimum Kolmogorov-Smirnov distance between the empirical distribution and the class of unimodal distributions. Larger values of the Dip-statistic signify departure from the null hypothesis of unimodality. P-values are calculated using the R-package of Maechler (2013). The p-value of this test is quite conservative.

INITIALIZE:

$K =$ number of clusters $= n$.

Sort data in ascending order and store them as: $\boldsymbol{x} = \{x_1, \ldots, x_n\}$.

Assign cluster mean $\{a_1, a_2, \ldots, a_n\}$ to them: $a_i = x_i$ for $i = 1, \ldots, n$.

Cluster size: $s_i = 1$, $i = 1, \ldots, n$.

Cluster Membership Indices of $\boldsymbol{x}$: $I(\boldsymbol{x}) = \{1, \ldots, n\}$.

WHILE $K > 1$:

Find the consecutive adjacent centroid distance standardized by cluster sizes:

$d(j, j+1) \leftarrow (a_{j+1} - a_j)/(s_j + s_{j+1})$

Find the clusters with minimum merging distance:

$j^\star \leftarrow \arg\min_{1 \le j \le K-1} d(j, j+1)$

Check if it is a Big Merge: $\min\{s_{j^\star}, s_{j^\star+1}\} > \lceil n\alpha \rceil$

IF Big Merge: Find and Store *Mass after merge* $= (s_{j^\star} + s_{j^\star+1})/n$ and

New Split = $\left\{ \max\{\boldsymbol{x}[I(\boldsymbol{x}) \text{ is } j^\star\}\, s_{j^\star} + \min\{\boldsymbol{x}[I(\boldsymbol{x}) \text{ is } (j^\star+1)]\}\, s_{j^\star+1} \right\} / (s_{j^\star} + s_{j^\star+1})$;

Merge the $j^\star$ and $(j^\star + 1)$ clusters and update the centroid and size of the new cluster:

$a_{j^\star} \leftarrow (s_{j^\star} a_{j^\star} + s_{j^\star+1} a_{j^\star+1})/(s_{j^\star} + s_{j^\star+1})$

$s_{j^\star} \leftarrow (s_{j^\star} + s_{j^\star+1})$

Reduce the number of clusters in path $K \leftarrow (K-1)$

Change cluster indices & cluster member indices of data according to the above reduction:

FOR k in $(j^\star + 1) : K$, $s_k \leftarrow s_{k+1}; a_k \leftarrow a_{k+1}$

FOR ALL I($\boldsymbol{x}$) $> j^\star$: reduce index by 1 ,i.e., $I(x) = I(x) - 1$

ADJUSTMENT: IF *Mass after merge* in the TOP SPLIT $< 50\%$, Stored.Splits = NULL;

OUTPUT `Stored Splits`.

Algorithm 3: α-thresholded BMT algorithm with truncation adjustment

E.2 Estimating the number of clusters

We compare the performance of the BMT with eight statistical methods that are popularly used for estimating the number of clusters in a dataset. A comparison study of 30 different approaches in Milligan and Cooper (1985) reports the approach in Caliński and Harabasz (1974) as being one of the best performing global methods. It prescribes maximizing the following index over k:

$$\text{CH}(k) = \frac{B(k)/(k-1)}{W(k)/(n-k)},$$

where $B(k)$ and $W(k)$ are respectively the between and the within clusters sum of squares for k clusters. Another popular approach, due to Krzanowski and Lai (1988), is based on the changes in the within clusters sum of squares as new clusters are formed, and seeks to maximize the following ratio over k:

$$\text{KL}(k) = \left| \frac{\text{DIFF}(k)}{\text{DIFF}(k+1)} \right| \text{ where DIFF}(k) = (k-1)^{2/p} W_{k-1} - k^{2/p} W_k.$$

Both these approaches are not defined for $k = 1$ and can not be used for testing population unimodality. Hartigan (1975) proposed using the smallest k for which the following ratio of the within cluster sum of squares is greater than 10:

$$\text{H}(k) = \left\{ \frac{W(k)}{W(k+1)} - 1 \right\} \Big/ (n-k-1).$$

It can be used for testing presence of only one cluster. Theoretical thresholds based on the F distribution can also be used. Gordon (1996) further sub-divides these approaches into local and global methods. Local methods consider individual pairs of clusters and check whether they should be merged. On the other hand, global methods incorporate the entire data in evaluating measures that are subsequently optimized as a function of the number of clusters. Note that the BMT is a local method. In addition to the above three methods, we compare BMT with the following five methods.

Given k clusters, for each data-point $\boldsymbol{x}_i$ the silhouette statistic of Kaufman and Rousseeuw (2009) uses: $\boldsymbol{a(i)}$ - the average distance of $\boldsymbol{x}_i$ to other points in its cluster, and $\boldsymbol{b(i)}$ - the average distance of $\boldsymbol{x}_i$ from points in its nearest neighboring cluster. It is given by $\text{sh}_k(i) = (b(i) - a(i))/\max\{a(i), b(i)\}$. Large values of $\text{sh}_k(i)$ signify good clustering.

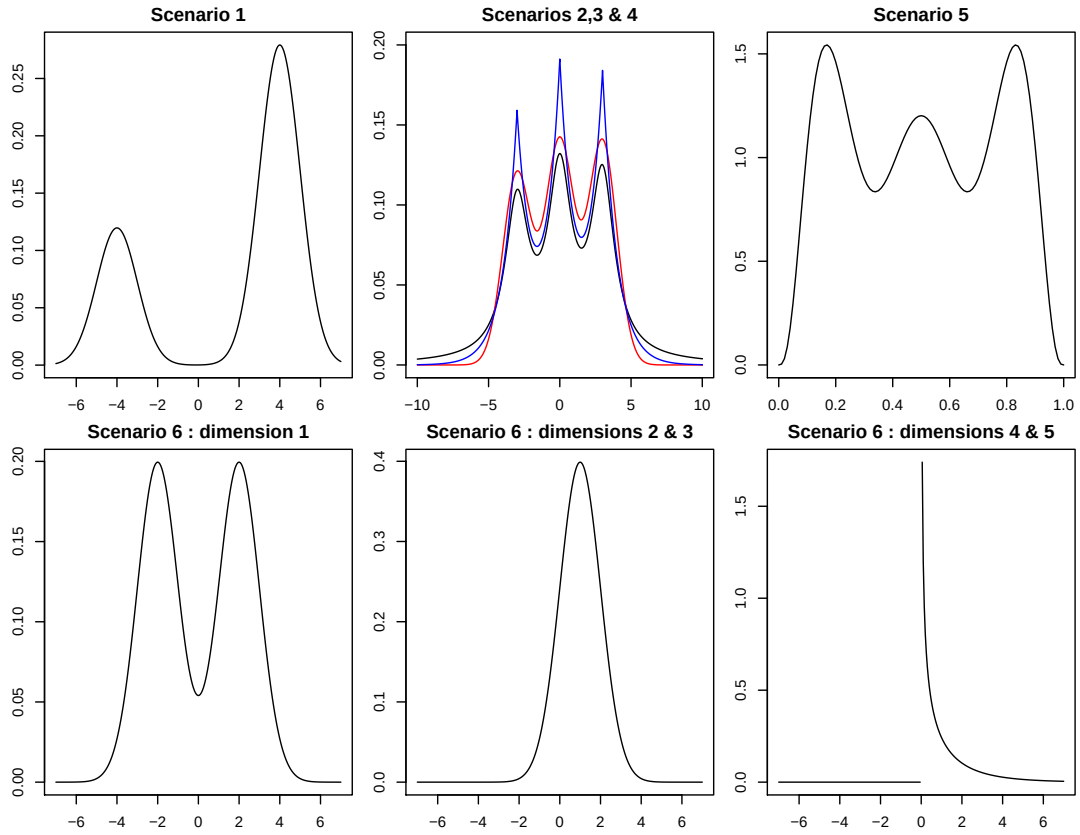

Figure S3: Plot of the univariate densities used in the different numerical experiments of Table 2 of the main paper.

A popular estimate of the optimal number of clusters is based on maximizing the average silhouette statistic,

$$\mathrm{KR}(k) = n^{-1} \sum_{i=1}^{n} \mathrm{sh}_k(i) \quad \text{over } k \geq 2.$$

The Gap statistic of Tibshirani *et al.* (2001) uses the 'elbow phenomenon' (Thorndike, 1953) by estimating the number of clusters at the transition point, where the decline in the within cluster dispersion first slackens. The goodness of clustering for k clusters is defined as:

$$\mathrm{Gap}_n(k) = \mathbb{E}_n^{\star}\{\log(\tilde{W}(k))\} - \log(\tilde{W}(k)),$$

where $\tilde{W}(k)$ is the size-normalized intra-cluster sums of squares. The expectation is over reference datasets and can be estimated by the mean of $\log \tilde{W}^*(k)$ over B i.i.d. datasets that are generated by sampling uniformly from the original dataset's range. The standard

deviation, std(k) of $\log \tilde{W}^*(k)$, is also recorded, and an estimate of the optimal number of clusters in the datasets is given by the smallest k for which the following holds:

$$\text{Gap}(k) \geq \text{Gap}(k+1) - (1 + B^{-1})^{-1/2}\text{std}(k+1).$$

Another method that uses the 'elbow phenomenon' is the Jump statistic of Sugar and James (2003). It again is a global method and involves computation of the mean square error (mse) associated with k clusters for different choices of k. The 'jumps' in the $p/2$ moments of the mean square errors are subsequently calculated: $J_k = \{\text{mse}(k)\}^{-p/2} - \{\text{mse}(k-1)\}^{-p/2}$ for $k = 1, \ldots, n$. The estimate of the number of clusters is the value of k that maximizes the jump J_k.

The prediction strength criterion of Tibshirani and Walther (2005) is computed by repeatedly dividing the data-set into two halves: train and test. For each value of $k \geq 1$, we cluster the test and training data into k groups and measure how well the training set cluster centroids predict co-memberships in the test set. For each of the k test clusters, the proportion of observation pairs that are also assigned to the same cluster by the training set centroids is computed. The prediction strength is the minimum of this value over the k test clusters. We have used the R-package of Hennig (2014) for computing the prediction strength criterion. 50 divisions of the data were used and the estimate of the number of clusters is the maximum the value of k which has prediction strength above 0.8.

The number of clusters estimate by Fang and Wang (2012) involves drawing 2 bootstrap samples from the data several times. For each of these bootstrap sample pairs and for each value of k, a global clustering instability measure (Wang, 2010), which reflects the distance between clusterings in paired samples, is computed. The number of clusters estimate minimizes the clustering instability aggregated over the bootstrap samples.

In Tables 2, 3 and 4 we report the performance of the aforementioned seven number of clusters estimation approaches (all except for the BMT and Jump), across the six simulation experiments, for the following clustering algorithms:

- Ward's algorithm (Ward Jr, 1963), which was implemented through the `NbCLust` package.
- Centroid-based clustering algorithm (Kaufman and Rousseeuw, 2009).
- PAM (Partitioning Around Medoids,) which is the most common realization of k-

medoid clustering algorithm (Kaufman and Rousseeuw, 1990). It uses a greedy search, which may not find the optimum solution but is faster than exhaustive search.

- CLARA (Clustering LARge Applications) algorithm of Kaufman and Rousseeuw (1986), which extends the k-medoids approach for large sample sizes. It first conducts a down-sampling and then clusters the down-sampled observations. Thereafter, it assigns all objects in the dataset to the clusters that were formed.

- GMM-merge, which conducts clustering by merging Gaussian mixture components from an initial mclust clustering (Hennig, 2010). The mixture components are merged in a hierarchical fashion. The merging criterion is computed for all pairs of current clusters and the two clusters with the extremal criterion value are merged. Then, criterion values are recomputed for the merged cluster and the merging process is continued up to the length dictated by number of clusters selection methods.

- GMM-BIC, which is based on Fraley and Raftery (2002). It fits a parameterized Gaussian mixture models by EM algorithm which is initialized by model-based hierarchical clustering.

Not all the above clustering algorithms were compatible with all of the number of clusters methods that we consider. Some of the combinations returned frequent computational errors due to non-convergence and related issues. We restricted ourselves to those pairs which were implementable in all our simulation set-ups. Also, as most of these clustering algorithms were very computationally expensive, we could not conduct 100 repetitions of the simulation scenarios as we had done for Table 2 in the main paper with the much faster k-means algorithm. We limited ourselves to 20 independent trails here, and most of the performance patterns were clearly visible by then. Figure S3, provided in this section, contains the plots of the densities used in our numerical experiments. In the second sub-plot of the top row we have the true population densities for the second, third and fourth simulation set-ups. In red, black and blue we have mixture of normals, t and double-exponential densities, respectively. Bootstrap stability with CLARA or with the GMM-merge algorithm and Prediction Strength with CLARA produced performances that are similar to those seen with the k-means clustering algorithm. The other combinations were not as effective.

True Population Density	Methods	Number of Clusters									
		1	2	3	4	5	6	7	8	9	10+
$0.3\,N(-4,1)+0.7\,N(4,1)$	CH	0	**3**	1	1	0	2	1	0	5	7
	KL	0	**11**	1	0	1	1	2	2	1	1
	Hartigan	0	**0**	7	4	6	3	0	0	0	0
	Silhouette	0	**20**	0	0	0	0	0	0	0	0
	Gap	0	**20**	0	0	0	0	0	0	0	0
$0.3\,N(-3,1)+0.35\,N(0,1)+0.35\,N(3,1)$	CH	0	1	**0**	0	0	2	0	1	5	11
	KL	0	2	**4**	1	1	2	1	6	1	2
	Hartigan	0	0	**11**	1	3	1	2	1	0	1
	Silhouette	0	9	**11**	0	0	0	0	0	0	0
	Gap	0	20	**0**	0	0	0	0	0	0	0
$0.3\,t_1(-3)+0.35\,t_1(0)+0.35\,t_1(3)$	CH	0	0	**0**	0	1	0	1	1	3	14
	KL	0	2	**5**	2	0	0	2	5	2	2
	Hartigan	0	0	**7**	4	1	0	3	2	1	2
	Silhouette	0	16	**4**	0	0	0	0	0	0	0
	Gap	0	10	**10**	0	0	0	0	0	0	0
$0.3\,\mathrm{dexp}(-3)+0.35\,\mathrm{dexp}(0)+0.35\,\mathrm{dexp}(3)$	CH	0	0	**3**	0	1	0	3	1	6	6
	KL	0	2	**3**	3	4	1	6	0	0	1
	Hartigan	0	0	**14**	3	0	0	1	1	0	1
	Silhouette	0	7	**13**	0	0	0	0	0	0	0
	Gap	0	19	**1**	0	0	0	0	0	0	0
$\{Beta(8,2)+Beta(5,5)+Beta(2,8)\}/3$	CH	0	0	**0**	0	0	0	2	2	6	10
	KL	0	4	**1**	5	4	3	0	1	0	2
	Hartigan	0	0	**11**	5	0	1	0	0	2	1
	Silhouette	0	11	**8**	1	0	0	0	0	0	0
	Gap	0	17	**3**	0	0	0	0	0	0	0
$\{0.5\,N(-2,1)+0.5\,N(2,1)\}$ $\otimes\, N(0,1)$ $\otimes\, N(0,1)$ $\otimes\, \chi^2_1$ $\otimes\, \chi^2_1$	CH	0	**19**	0	1	0	0	0	0	0	0
	KL	0	**5**	3	3	1	2	2	0	3	1
	Hartigan	0	**0**	5	13	0	2	0	0	0	0
	Silhouette	0	**9**	5	6	0	0	0	0	0	0
	Gap	0	**20**	0	0	0	0	0	0	0	0

Table 2: Number of clusters detected in 20 trials for the Ward's clustering algorithm

True Population Density	Methods	Number of Clusters									
		1	2	3	4	5	6	7	8	9	10+
$0.3\,N(-4,1)+0.7\,N(4,1)$	CH	0	**17**	0	0	0	0	0	1	1	1
	KL	0	**1**	0	1	2	1	1	1	9	4
	Hartigan	0	**0**	0	1	2	4	0	4	7	2
	Silhouette	0	**20**	0	0	0	0	0	0	0	0
	Gap	0	**20**	0	0	0	0	0	0	0	0
$0.3\,N(-3,1)+0.35\,N(0,1)+0.35\,N(3,1)$	CH	0	1	**1**	0	0	0	1	3	6	8
	KL	0	2	**1**	4	2	2	0	4	3	2
	Hartigan	0	0	**12**	2	2	1	2	1	0	0
	Silhouette	0	5	**3**	8	3	1	0	0	0	0
	Gap	0	17	**3**	0	0	0	0	0	0	0
$0.3\,t_1(-3)+0.35\,t_1(0)+0.35\,t_1(3)$	CH	0	1	**1**	0	0	5	3	1	2	7
	KL	0	3	**4**	3	5	2	1	2	0	0
	Hartigan	0	0	**12**	4	2	0	0	0	0	2
	Silhouette	0	20	**0**	0	0	0	0	0	0	0
	Gap	0	6	**14**	0	0	0	0	0	0	0
$0.3\,\mathrm{dexp}(-3)+0.35\,\mathrm{dexp}(0)+0.35\,\mathrm{dexp}(3)$	CH	0	0	**0**	0	1	2	1	2	5	9
	KL	0	1	**4**	2	2	3	1	4	2	1
	Hartigan	0	0	**8**	2	5	4	0	0	1	0
	Silhouette	0	11	**0**	0	3	3	3	0	0	0
	Gap	0	19	**1**	0	0	0	0	0	0	0
$\{Beta(8,2)+Beta(5,5)+Beta(2,8)\}/3$	CH	0	0	**1**	0	0	0	2	1	5	11
	KL	0	2	**1**	5	1	1	3	0	5	2
	Hartigan	0	0	**11**	6	2	1	0	0	0	0
	Silhouette	0	5	**12**	3	0	0	0	0	0	0
	Gap	0	19	**1**	0	0	0	0	0	0	0
$\{0.5\,N(-2,1)+0.5\,N(2,1)\}$ $\otimes\, N(0,1)$ $\otimes\, N(0,1)$ $\otimes\, \chi^2_1$ $\otimes\, \chi^2_1$	CH	0	**7**	4	3	3	1	0	1	1	0
	KL	0	**5**	4	2	1	0	3	2	1	2
	Hartigan	0	**0**	4	5	2	2	3	0	2	2
	Silhouette	0	**20**	0	0	0	0	0	0	0	0
	Gap	0	**20**	0	0	0	0	0	0	0	0

Table 3: Number of clusters detected in 20 trials for the centroid clustering algorithm

True Population Density	Methods	Number of Clusters									
		1	2	3	4	5	6	7	8	9	10+
$0.3\,N(-4,1) + 0.7\,N(4,1)$	BS and GMM-BIC	0	**20**	0	0	0	0	0	0	0	0
	BS and GMM-merge	0	**20**	0	0	0	0	0	0	0	0
	BS and CLARA	0	**20**	0	0	0	0	0	0	0	0
	PS and CLARA	0	**20**	0	0	0	0	0	0	0	0
	BS and PAM	0	**20**	0	0	0	0	0	0	0	0
	PS and PAM	0	**6**	14	0	0	0	0	0	0	0
$0.3\,N(-3,1) + 0.35\,N(0,1) + 0.35\,N(3,1)$	BS and GMM-BIC	0	0	**19**	0	0	0	0	0	0	1
	BS and GMM-merge	0	0	**20**	0	0	0	0	0	0	0
	BS and CLARA	0	1	**19**	0	0	0	0	0	0	0
	PS and CLARA	0	0	**20**	0	0	0	0	0	0	0
	BS and PAM	0	2	**18**	0	0	0	0	0	0	0
	PS and PAM	5	15	**0**	0	0	0	0	0	0	0
$0.3\,t_1(-3) + 0.35\,t_1(0) + 0.35\,t_1(3)$	BS and GMM-BIC	0	17	**3**	0	0	0	0	0	0	0
	BS and GMM-merge	0	20	**0**	0	0	0	0	0	0	0
	BS and CLARA	0	1	**0**	3	8	3	4	1	0	0
	PS and CLARA	6	7	**2**	4	1	0	0	0	0	0
	BS and PAM	0	3	**0**	0	0	2	3	6	6	0
	PS and PAM	20	0	**0**	0	0	0	0	0	0	0
$0.3\,\mathrm{dexp}(-3) + 0.35\,\mathrm{dexp}(0) + 0.35\,\mathrm{dexp}(3)$	BS and GMM-BIC	0	1	**5**	0	4	5	2	0	3	0
	BS and GMM-merge	0	0	**20**	0	0	0	0	0	0	0
	BS and CLARA	0	0	**20**	0	0	0	0	0	0	0
	PS and CLARA	0	0	**20**	0	0	0	0	0	0	0
	BS and PAM	0	0	**20**	0	0	0	0	0	0	0
	PS and PAM	11	9	**0**	0	0	0	0	0	0	0
$\big\{Beta(8,2) + Beta(5,5) + Beta(2,8)\big\}/3$	BS and GMM-BIC	0	0	**1**	2	1	5	2	3	1	5
	BS and GMM-merge	0	20	**0**	0	0	0	0	0	0	0
	BS and CLARA	0	12	**5**	0	1	0	0	0	2	0
	PS and CLARA	1	9	**9**	1	0	0	0	0	0	0
	BS and PAM	0	7	**12**	0	1	0	0	0	0	0
	PS and PAM	20	0	**0**	0	0	0	0	0	0	0
$\{0.5\,N(-2,1) + 0.5\,N(2,1)\}$ $\otimes\ N(0,1)\ \otimes\ N(0,1)$ $\otimes\ \chi^2_1\ \otimes\ \chi^2_1$	BS and GMM-BIC	0	**0**	0	0	0	0	0	0	0	20
	BS and GMM-merge	0	**20**	0	0	0	0	0	0	0	0
	BS and CLARA	0	**20**	0	0	0	0	0	0	0	0
	PS and CLARA	0	**20**	0	0	0	0	0	0	0	0
	BS and PAM	0	**5**	0	0	0	0	0	1	0	14
	PS and PAM	15	**5**	0	0	0	0	0	0	0	0

Table 4: Number of clusters detected in 20 trials by the prediction strength (PS) and bootstrap stability (BS) methods for selecting the number of clusters. Four clustering algorithms, CLARA, PAM, GMM-BIC and GMM-merge, were used in combination with each of the two aforementioned number of clusters methods whenever the two parts were compatible in the R-package *fpc*.

E.3 Performance on Large Data Sets

In Table 5 we report the performance of the BMT across 4 different simulation examples involving large samples generated from Gaussian mixtures. For each example, 100 independent data sets are generated from the population density and the distribution of the number of clusters detected by the BMT is reported (the frequencies are in parenthesis). We also report the average and the standard deviation of the Mean Square Error (MSE) over the cases where the the true number of clusters is correctly detected. The oracle MSE is calculated based on the partition that uses the minima of the true population density.

Population Density	Sample Size	Time in Sec per Replicate	Number of Clusters	MSE		Oracle MSE
				Mean	SD	
$\{N(-2.5,1)+N(0,1)+N(2.5,1)\}/3$	10^4	0.70	2 (3), **3 (96)**, 4 (1)	0.6817	0.0405	0.6564
$0.5\,N(-5,1)+0.25\,N(0,1)+0.25\,N(5,1)$	5×10^4	11.58	**3 (100)**	1.1152	0.0125	0.9769
$\{N(-1.1,1)+N(1.1,1)\}/2$	10^5	44.07	1 (25), **2 (70)**, 3 (5)	0.6905	0.0258	0.6789
$\{N(0,1)\pm N(4,1)\pm N(8,1)\}/5$	10^5	38.98	**5 (100)**	0.8909	0.0036	0.8909

Table 5: Performance of the BMT on simulated datasets of large sample sizes

Overall, the BMT correctly detects the true number of clusters with high certainly. Also, the average MSE is observed to be very close to the Oracle one. To demonstrate the scalability of the proposed method, we report the average elapsed time (in seconds) per replication. The numerical experiments were performed at the Center for High Performance Computing (`http://hpcc.usc.edu`) of the University of Southern California. The computations were done in R version 3.1.1 on Dual Quadcore Intel Xeon 2.33 GHz, 16GB Memory nodes. We used the `Snowfall` package of Knaus (2013) to distribute computations over 100 CPUs.

Most of the competitor methods discussed above fail to accommodate the large sample sizes of the data sets in Table 5. To get a better understanding of the relative scalibility of the BMT, we implemented the HMAC algorithm of Li *et al.* 2007, which can handle large sample sizes. HMAC is a non-parametric mode identification based clustering approach, which produces clusters in hierarchical levels. Compared to HMAC, the BMT was found to be substantially faster in producing the entire hierarchical path. On an iMac desktop with 2.9 GHz Intel Core i5 processor and 8 GB memory, for grids of 20 hierarchical levels, the

R-code associated with the HMAC paper required approximate run times of 3 minutes, 70 minutes and 4 hours for the sample sizes of 10000, 50000 and 100000, respectively. On the other hand, the corresponding average run times for the BMT were 10 seconds, 5 minutes and 20 minutes, respectively.

E.4 BMT & Sub-population Analysis in Single Cell Virology

We demonstrate an application of our clustering method in an immunology study conducted at single cell level. Emerging technologies (Wang and Bodivitz, 2011) have recently enabled us to collect proteomic data sets at single cell resolution. These data sets reflect the variations of protein expressions across cells and need clustering techniques for detection of cellular sub-populations. Typically, sub-populations are detected by core-protein expressions based cluster analysis of the samples, and the signaling expressions of the resultant sub-populations are subsequently studied. In Figures S4 and S5 we display the results of the BMT induced clustering on the virology datasets of Sen *et al.* (2014). Figure S6 shows the post-clustering, sub-population level signaling expressions.

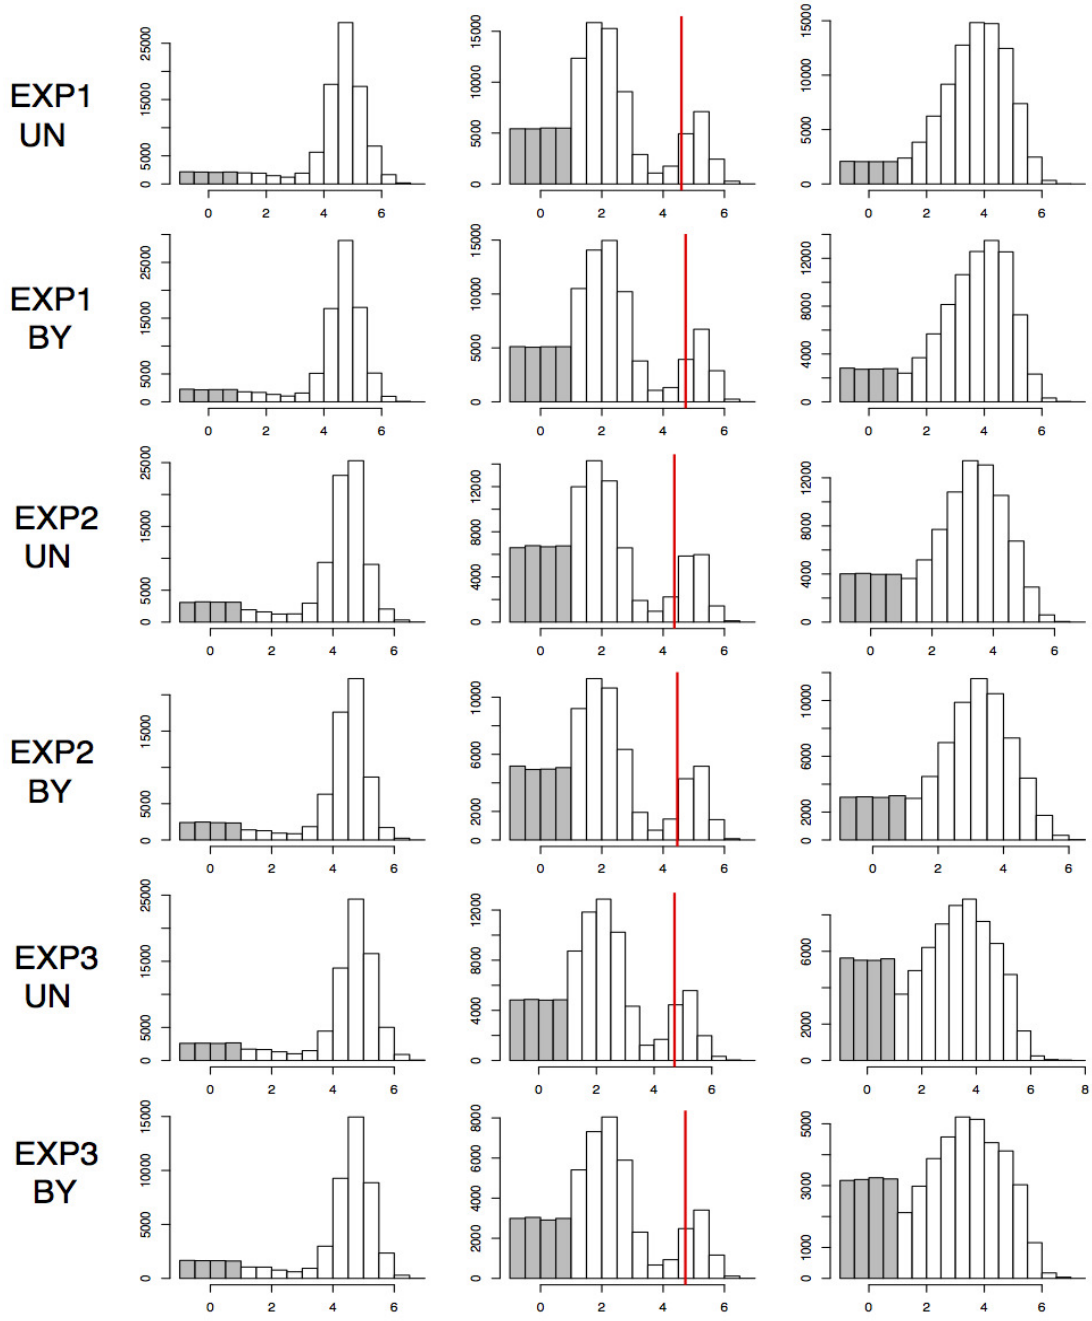

Figure S4: Across columns we have histograms of the expression values of the proteins CD4, CD8 and CD45RA, respectively. Along rows, from top to bottom, we have the histograms of the Uninfected and Bystander population, respectively, for the independent experiments I-III. The shaded gray region denotes unexpressed values. Splits in the expression values (if any) detected by BMT are shown by vertical red lines.

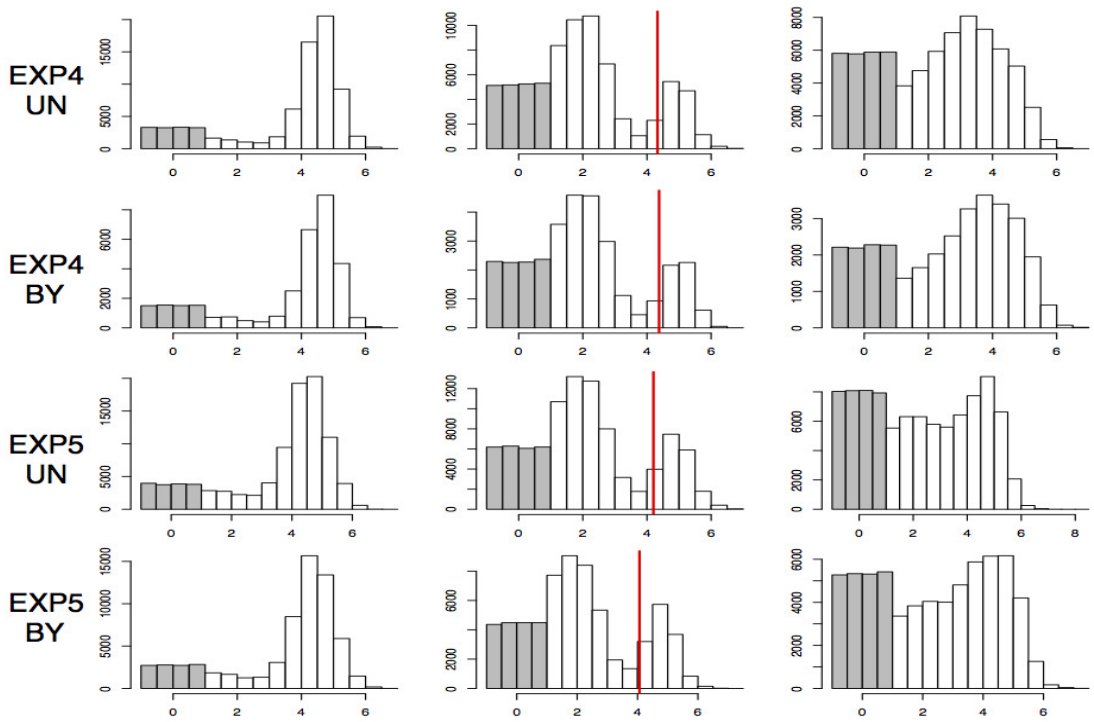

Figure S5: Across columns we have histograms of the expression values of CD4, CD8 and CD45RA. Across rows are the histograms of UN and BY populations for Experiments IV-V. Splits detected by BMT (if any) are shown by red lines.

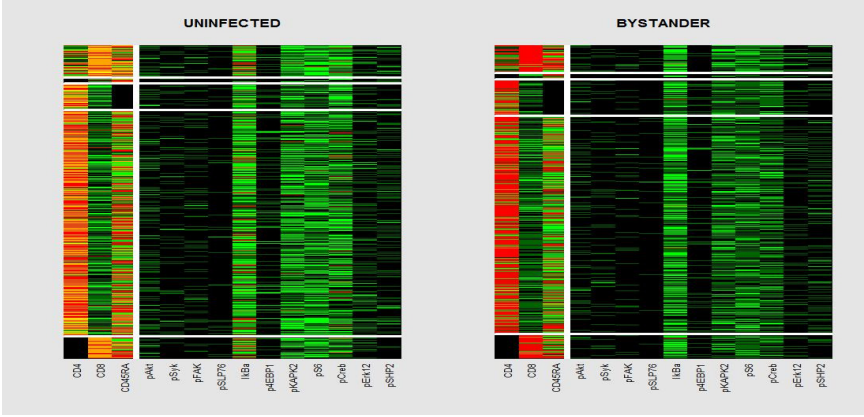

Figure S6: The above plot shows the heatmaps of the protein expression values (in order of decreasing intensity: Red, Yellow, Green and Black) of the Uninfected and Bystander populations in Experiment I. The horizontal white lines demarcate the five major sub-population detected by BMT algorithm, based on the expression of the three surface markers on the left of the vertical white line. The proteins on the right of the vertical line are associated with cell-signaling. The heatmaps are standardized separately for the two populations.

References

Caliński, T. and Harabasz, J. (1974). A dendrite method for cluster analysis. *Communications in Statistics-theory and Methods* **3**, 1, 1–27.

Fang, Y. and Wang, J. (2012). Selection of the number of clusters via the bootstrap method. *Computational Statistics & Data Analysis* **56**, 3, 468–477.

Fraley, C. and Raftery, A. E. (2002). Model-based clustering, discriminant analysis, and density estimation. *Journal of the American statistical Association* **97**, 458, 611–631.

Gordon, A. D. (1996). Null models in cluster validation. In *From data to knowledge*, 32–44. Springer.

Hall, P. and York, M. (2001). On the calibration of silverman's test for multimodality. *Statistica Sinica* **11**, 2, 515–536.

Hartigan, J. A. (1975). *Clustering algorithms*. Wiley.

Hartigan, J. A. and Hartigan, P. (1985). The dip test of unimodality. *The Annals of Statistics* 70–84.

Hennig, C. (2010). Methods for merging gaussian mixture components. *Advances in data analysis and classification* **4**, 1, 3–34.

Hennig, C. (2014). *fpc: Flexible procedures for clustering*. R package version 2.1-9.

Kaufman, L. and Rousseeuw, P. (1986). *Clustering large data sets*. Elsevier.

Kaufman, L. and Rousseeuw, P. J. (1990). Partitioning around medoids (program pam). *Finding groups in data: an introduction to cluster analysis* 68–125.

Kaufman, L. and Rousseeuw, P. J. (2009). *Finding groups in data: an introduction to cluster analysis*, vol. 344. John Wiley & Sons.

Kim, J. and Pollard, D. (1990). Cube root asymptotics. *The Annals of Statistics* 191–219.

Knaus, J. (2013). *snowfall: Easier cluster computing (based on snow).* R package version 1.84-6.

Krzanowski, W. J. and Lai, Y. (1988). A criterion for determining the number of groups in a data set using sum-of-squares clustering. *Biometrics* 23–34.

Li, J., Ray, S., and Lindsay, B. G. (2007). A nonparametric statistical approach to clustering via mode identification. *Journal of Machine Learning Research* **8**, 8, 1687–1723.

Maechler, M. (2013). *diptest: Hartigan's dip test statistic for unimodality - corrected code*. R package version 0.75-5.

Milligan, G. W. and Cooper, M. C. (1985). An examination of procedures for determining the number of clusters in a data set. *Psychometrika* **50**, 2, 159–179.

Sen, N., Mukherjee, G., Sen, A., Bendall, S., Sung, P., Nolan, G., and Arvin, A. (2014). Single-cell mass cytometry analysis of human tonsil t cell remodeling by varicella zoster virus. *Cell Reports* **8**, 2, 633 – 645.

Silverman, B. W. (1981). Using kernel density estimates to investigate multimodality. *Journal of the Royal Statistical Society. Series B (Methodological)* 97–99.

Sugar, C. A. and James, G. M. (2003). Finding the number of clusters in a dataset. *Journal of the American Statistical Association* **98**, 463.

Thorndike, R. L. (1953). Who belongs in the family? *Psychometrika* **18**, 4, 267–276.

Tibshirani, R. and Walther, G. (2005). Cluster validation by prediction strength. *Journal of Computational and Graphical Statistics* **14**, 3, 511–528.

Tibshirani, R., Walther, G., and Hastie, T. (2001). Estimating the number of clusters in a data set via the gap statistic. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **63**, 2, 411–423.

van der Vaart, A. W. (1998). *Asymptotic Statistics*. Cambridge University Press.

Vollmer, S., Holzmann, H., and Schwaiger, F. (2013). Peaks vs components. *Review of Development Economics* **17**, 2, 352–364.

Wang, D. J. and Bodivitz, S. (2011). Single cell analysis: the new frontier in omics. *Nature Methods* .

Wang, J. (2010). Consistent selection of the number of clusters via crossvalidation. *Biometrika* **97**, 4, 893–904.

Ward Jr, J. H. (1963). Hierarchical grouping to optimize an objective function. *Journal of the American statistical association* **58**, 301, 236–244.